

CAMBRIDGE TRACTS IN MATHEMATICS

175

**THE LARGE SIEVE AND
ITS APPLICATIONS**
ARITHMETIC GEOMETRY, RANDOM
WALKS AND DISCRETE GROUPS

EMMANUEL KOWALSKI



CAMBRIDGE UNIVERSITY PRESS

This page intentionally left blank

CAMBRIDGE TRACTS IN MATHEMATICS

General Editors

B. BOLLOBÁS, W. FULTON, A. KATOK, F. KIRWAN,
P. SARNAK, B. SIMON, B. TOTARO

**175 The Large Sieve and its Applications:
Arithmetic Geometry, Random Walks and Discrete Groups**

The Large Sieve and its Applications

Arithmetic Geometry, Random Walks and
Discrete Groups

E. KOWALSKI

Swiss Federal Institute of Technology (ETH), Zürich



CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS

Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo

Cambridge University Press

The Edinburgh Building, Cambridge CB2 8RU, UK

Published in the United States of America by Cambridge University Press, New York

www.cambridge.org

Information on this title: www.cambridge.org/9780521888516

© E. Kowalski 2008

This publication is in copyright. Subject to statutory exception and to the provision of relevant collective licensing agreements, no reproduction of any part may take place without the written permission of Cambridge University Press.

First published in print format 2008

ISBN-13 978-0-511-39887-2 eBook (EBL)

ISBN-13 978-0-521-88851-6 hardback

Cambridge University Press has no responsibility for the persistence or accuracy of urls for external or third-party internet websites referred to in this publication, and does not guarantee that any content on such websites is, or will remain, accurate or appropriate.

Pour les soixante ans de Jean–Marc Deshouillers

Contents

<i>Preface</i>	<i>page xi</i>
<i>Acknowledgments</i>	<i>xvi</i>
<i>Prerequisites and notation</i>	<i>xvii</i>
1 Introduction	1
1.1 Presentation	1
1.2 Some new applications of the large sieve	4
2 The principle of the large sieve	8
2.1 Notation and terminology	8
2.2 The large sieve inequality	9
2.3 Duality and ‘exponential sums’	18
2.4 The dual sieve	22
2.5 General comments on the large sieve inequality	25
3 Group and conjugacy sieves	32
3.1 Conjugacy sieves	32
3.2 Group sieves	34
3.3 Coset sieves	36
3.4 Exponential sums and equidistribution for group sieves	40
3.5 Self-contained statements	42
4 Elementary and classical examples	45
4.1 The inclusion-exclusion principle	45
4.2 The classical large sieve	48
4.3 The multiplicative large sieve inequality	57
4.4 The elliptic sieve	59
4.5 Other examples	67

5	Degrees of representations of finite groups	70
5.1	Introduction	70
5.2	Groups of Lie type with connected centres	72
5.3	Examples	82
5.4	Some groups with disconnected centres	83
6	Probabilistic sieves	87
6.1	Probabilistic sieves with integers	87
6.2	Some properties of random finitely presented groups	94
7	Sieving in discrete groups	101
7.1	Introduction	101
7.2	Random walks in discrete groups with Property (τ)	105
7.3	Applications to arithmetic groups	113
7.4	The cases of $SL(2)$ and $Sp(4)$	119
7.5	Arithmetic applications	127
7.6	Geometric applications	132
7.7	Explicit bounds and arithmetic transitions	145
7.8	Other groups	151
8	Sieving for Frobenius over finite fields	154
8.1	A problem about zeta functions of curves over finite fields	155
8.2	The formal setting of the sieve for Frobenius	160
8.3	Bounds for sieve exponential sums	164
8.4	Estimates for sums of Betti numbers	168
8.5	Bounds for the large sieve constants	171
8.6	Application to Chavdarov's problem	175
8.7	Remarks on monodromy groups	187
8.8	A last application	193
Appendix A	Small sieves	197
A.1	General results	197
A.2	An application	201
Appendix B	Local density computations over finite fields	204
B.1	Density of cycle types for polynomials over finite fields	204
B.2	Some matrix densities over finite fields	210
B.3	Other techniques	218

Appendix C	Representation theory	220
C.1	Definitions	220
C.2	Harmonic analysis	223
C.3	One-dimensional representations	226
C.4	The character tables of $GL(2, \mathbf{F}_q)$ and $SL(2, \mathbf{F}_q)$	227
Appendix D	Property (T) and Property (τ)	232
D.1	Property (T)	232
D.2	Properties and examples	233
D.3	Property (τ)	236
D.4	Shalom's theorem	238
Appendix E	Linear algebraic groups	245
E.1	Basic terminology	245
E.2	Galois groups of characteristic polynomials	249
Appendix F	Probability theory and random walks	254
F.1	Terminology	254
F.2	The Central Limit Theorem	257
F.3	The Borel–Cantelli lemmas	258
F.4	Random walks	259
Appendix G	Sums of multiplicative functions	262
G.1	Some basic theorems	262
G.2	An example	264
Appendix H	Topology	268
H.1	The fundamental group	268
H.2	Homology	275
H.3	The mapping class group of surfaces	276
<i>References</i>		283
<i>Index</i>		289

Preface

‘The Romans,’ Roger and the Reverend Dr. Paul de la Nuit were drunk together one night, or the vicar was, ‘the ancient Roman priests laid a sieve in the road, and then waited to see which stalks of grass would come up through the holes.’

Thomas Pynchon, ‘Gravity’s Rainbow’

These notes arose, by the long and convoluted process that research often turns out to be, from a supposedly short addition to my paper [80]. This is a story that is certainly typical of much of scientific research, and since I always find this fascinating, and hardly visible from the outside once a paper or book is published,¹ I will summarize the events briefly. Readers who like science rather dry or dour may wish to start reading Chapter 1.

The original ambition was simply to extend the large sieve bound for Frobenius conjugacy classes of this first paper to the stronger form classically due to Montgomery, which would mean that ‘small sieve’ applications would become possible. The possibility of this extension seemed clear to me, as well as the relative paucity of new applications.² At the same time, it seemed natural to ‘axiomatize’ the setting in a way allowing an identical treatment of the classical large sieve inequality and this newer variant, and this seemed a worthwhile enough goal.

All this should not have taken very long, either in time or space, except that inevitable delays due to teaching and other duties led to the thought that maybe other applications of this abstract form of sieve would be possible, and could be

¹ A striking recent instance of this process is described by A. Wiles in the introduction to his paper proving Fermat’s Great Theorem.

² In large sieve situations, applying the best small sieve bound gives very small gains, whereas small sieve cases, by definition, can be handled by small sieves, which were already sufficiently general to handle the ‘obvious’ applications, and in fact strong enough to prove lower bounds in some contexts.

briefly discussed in the course of the paper, which would thus become stronger. A natural fit, given my background and the emphasis on random matrices as an interpretation of the results of [80], was to think of trying to prove, e.g., that a ‘generic’ unimodular integral $n \times n$ matrix has irreducible characteristic polynomial (or maximal splitting field), as an application of the large sieve applied to $SL(n, \mathbf{Z})$. I started thinking about this problem, seeing clearly that harmonic analysis of automorphic forms on $SL(n, \mathbf{Z}) \backslash SL(n, \mathbf{R})$ would be called for, and that this would require some learning on my part for $n \geq 3$. Clearly this would be material for *another* paper, a quite interesting one since I knew of no previous use of sieve in such situations. Because of the strong link to spectral theory of automorphic forms, I was pretty sure I would have heard of it if published papers on this topic existed; as it was, there were results of Duke, Rudnick and Sarnak [33] (and their later extensions) giving asymptotic formulas for the *number* of unimodular matrices with bounded norm, but not for the more general ‘exponential sums’ arising from the sieve theory.

In the meantime, D. Zywinia sent me his preprint (*‘The large sieve and Galois representations’*, 2007) which contained a slightly different formulation of an abstract form of the large sieve, with applications to distribution of Frobenius elements in number fields, specifically to the Lang–Trotter Conjecture. His sieve axioms were in many respects more general than mine, except for one condition which I had to introduce in [80] because of specific features of the arithmetic of varieties over finite fields (the difference between arithmetic and geometric fundamental groups). Still, where his conditions were more general, I could in fact very easily assume the same generality, and reading his preprint led me to rewrite mine in this light. This did not bring new applications. On the other hand, as I was reading (mostly for the pleasure of it) the nice book by P. de la Harpe on geometric group theory [57], I thought that one could also try to use as targets of sieves the subsets of groups defined by word length (with respect to some system of generators) being smaller than some quantity. However, not knowing much about this topic, this was mostly speculative.

But around the same time, I. Rivin posted a preprint [108] on arXiv (www.arXiv.org) which directly mentioned the problem of irreducibility of characteristic polynomials of unimodular matrices. He also mentioned the results of Duke, Rudnick and Sarnak but did not prove that ‘most’ matrices have this property. What he managed to prove was an analogue of the more combinatorial variant: instead of looking at balls in the word-length metric, rather he was looking at random walks on the group of length $k \rightarrow +\infty$. His method for detecting irreducibility was similar to the ‘old’ method used by van der Waerden for integral polynomials with bounded height, combined with results of Chavdarov [22] (which already played a role in [80], one of the results

of which was indeed a strong quantitative strengthening of Chavdarov's main result, following Gallagher's large sieve strengthening [46] of van der Waerden's result), and in particular the statement proved was qualitative and did not give explicit bounds for the probability of having a reducible characteristic polynomial.

A remarkable novel feature of Rivin's work was the new applications he discussed, which concerned 'generic' properties of automorphisms either of compact connected surfaces or free groups. In each case, the action of such elements on a free abelian group (the homology of the surface or abelianization of the free group, respectively) was sufficient to detect an interesting condition by looking at the corresponding characteristic polynomial. Rivin thus proved in a very simple way a (special case of a) result of Maher [96]: the probability that the k -th step of a random walk on the mapping class group of a surface of genus g is pseudo-Anosov tends to 1 as $k \rightarrow +\infty$.

As I mentioned to Rivin that I had been working with the large sieve with applications to characteristic polynomials in mind, he told me that Bourgain, Gamburd and Sarnak were investigating issues related to sieve in arithmetic groups and forwarded their preprint [14]. This work was, in small sieve contexts, concerned with showing that orbits of certain subgroups G of arithmetic groups acting on \mathbf{Z}^n contain infinitely many points with prime (or almost prime) coordinates. What was clearly explained was that, apart from fairly standard sieve machinery going back to Brun or Selberg, the crucial feature that must be exploited (and proved) is the expanding property of congruence quotients of the group G .

As I became aware of these very interesting developments, my paper remained unchanged. Or rather, what was expanding in it was a 'sidebar' having to do with natural questions suggested by the sieve framework: what is the largest dimension of an irreducible representation of a finite group of Lie type, such as $SL(n, \mathbf{Z}/\ell\mathbf{Z})$ or $Sp(2g, \mathbf{Z}/\ell\mathbf{Z})$, and what is the sum of those dimensions? This had already puzzled me while writing [80], where I used 'trivial' bounds for those quantities. As I tried once more to get some understanding of the theory of Deligne–Lusztig characters which describes the representations of such groups, I finally wrote to F. Digne and J. Michel, with the feeling that this must certainly be known, but hidden somewhere inaccessible to 'simple' searches in *Mathematical Reviews*. However, J. Michel did not know if the first question had been considered (he pointed out the papers of Gow [50] and Vinroot [129] concerning the second problem). Based on his indications, I managed to write down a proof of the estimate which I had found 'reasonable' to expect.

Finally summer vacation came. Then, in a short time, I found and wrote down a new amusing application of the sieve to the study of denominators of rational points on elliptic curves, which was a good example of the ‘abstract’ framework. More importantly, Rivin’s use of random walks prompted me to generalize the sieve context to that of estimating the measure of some ‘sifted set’, and not necessarily its cardinality, in order to incorporate applications having to do with general random walks. And using Property (τ) for discrete groups together with some nice probabilistic ideas described in the survey on random walks on groups by L. Saloff-Coste [111], I obtained an effective form of Rivin’s irreducibility theorem for random walks on $SL(n, \mathbf{Z})$ or $Sp(2g, \mathbf{Z})$.

At this point, I felt that I merely needed to polish a few things and then send the paper to a well-chosen journal. I was wondering if splitting it into multiple parts might not be better (something I usually strongly dislike), since its growing mathematical spread, while appealing, obviously made it difficult to find a single referee: by this time, the crucial insights were from analytic number theory, the tools ranged from representation theory, including Deligne–Lusztig theory, to Property (τ) and the Riemann Hypothesis over finite fields, not to mention the use of probabilistic vocabulary. And familiarity with [80] was quite obviously assumed . . .

But then I realized that the very basic formal part of the large sieve was unduly complicated and framed in the wrong way, bonding the method with group theory much too early (the title at the time was ‘The algebraic principle of the large sieve’, a joking pun on [98]). By moving the group theory to a different part of the argument (the choice of a suitable orthonormal basis for finite-dimensional Hilbert spaces), the principle of the sieve could be both simplified and generalized once more. In retrospect, nothing seems more obvious, but the simpler form had been completely obscured by the force of habit together with the fact that all applications I knew were linked with a group and its representation theory.

So I rewrote much of the beginning part and adapted the rest; by this time the paper was around 55 (full) pages long. After some more hesitation, some more feature-creep, and taking advice from P. Sarnak and A. Granville, getting this text in a journal seemed less and less practical. Because of the many applications, I wanted the paper to be accessible to as large an audience as possible, and the style of the writing appeared to me to become unsuitable for, say, geometers interested in the stronger form of Maher’s and Rivin’s results (I had realized, looking at [96] quite late, that my bound for characteristic polynomials of $Sp(2g, \mathbf{Z})$ implied a solution to a further question of Maher, namely the *transience* of the set of non-pseudo-Anosov elements during a random walk on the mapping class group).

The outcome of this process is that I have expanded the paper to a short book, adding brief surveys of most of the important material that may not be known to all readers. This includes the representation theory of finite groups, Property (τ) (and Property (T)) – with a sketch of the proof of Property (T) for $SL(n, \mathbf{Z})$ due to Shalom [124], sums of multiplicative functions, probability theory and random walks, and the mapping class groups of surfaces. Of course, for some of these, I have no claim to expertise and the surveys should only be thought of as delineating the basic definitions and some basic information which I found especially interesting (or beautiful!) while learning about the subject.

All this will, I hope, have both the effect of making the text readable for non-analytic number theorists that may have potential use of ideas related to the large sieve, and to make analytic number theorists aware of some potential areas where their ideas might be useful.

Acknowledgments

As the preface shows, a lot of people have had a great influence on the final appearance of this work beyond the impetus of [80]. I mention again in particular D. Zywin, who developed an abstract setup of the large sieve similar to the conjugacy sieve described in Chapter 3, which prompted me on more than one occasion to evolve my own version; and I. Rivin, whose work suggested the probabilistic sieve setting, and who also mentioned to me the work of Bourgain, Sarnak and Gamburd. I also wish to thank P. Sarnak for sending me a copy of his email to his coauthors. Finally, I thank J. Michel for providing the ideas of the proof of Proposition 5.5 and explaining some basic properties of representations of finite groups of Lie type; M. Burger for information concerning Property (T) and Property (τ) , and for correcting my misunderstanding of his paper [18]; P. Duchon and M.-L. Chabanol for help, advice and references concerning probability theory and graph theory; K. Belabas for suggestions concerning numerical experiments; D. Khoshnevisan for providing a correct proof of one probabilistic statement; and J. Wu for explaining some points concerning [88]. Also, I wish to thank F. Jouve for finding many small mistakes and imprecisions in the original drafts.

Work on this book was partially supported by the ANR (L'Agence Nationale de la Recherche) Project ARITHMATICS. Some preliminary results were presented during the conference organized by this project in Bordeaux in April 2006, and the remarks of participants were very helpful in shaping the later evolution of the ideas presented here. A much shorter preliminary version of this book was also posted on arXiv as `arXiv:math.NT/0610021`.

Prerequisites and notation

There are two types of readers for whom this book is written: some who are knowledgeable about analytic number theory, and maybe very familiar with sieve methods, and who (we hope) will find the new and unfamiliar applications of interest; and some who are interested in a specific application (e.g., those around properties of mapping class groups, or zeta functions of algebraic varieties over finite fields, or random walks on discrete groups), but not necessarily in all of them, and who may not be familiar with the principles of analytic number theory.

Fortunately, there is in fact very little prerequisite for most of the book; the basic principle of the large sieve uses nothing more than basic linear algebra and analysis (finite-dimensional Hilbert spaces). When it comes to applications, where more sophisticated tools are often involved, we follow the policy of defining from scratch all notions that appear, and provide the reader with precise references for all facts we use about such topics as elliptic curves, discrete groups, algebraic groups, random walks and harmonic analysis. The only (partial) exception is in Chapter 8 where we need the machinery of ℓ -adic sheaves over finite fields, and their cohomology. But even then, the statements of the applications of the sieve (at least) should be understandable by any reader, and we hope that the mechanism of the proofs is explained clearly enough that analytic number theorists will be able to benefit from reading this chapter.

We now summarize the most common notation. Less standard notation will be explained in each chapter when first used (see in particular the beginning of Chapter 2), and moreover the appendices contain quick surveys of the definitions of (almost) all mathematical terms which occur in the book.

As usual, $|X|$ denotes the cardinality of a set; however if X is a measure space with measure μ , we sometimes write $|X|$ instead of $\mu(X)$.

By $f \ll g$ for $x \in X$, or $f = O(g)$ for $x \in X$, where X is an arbitrary set on which f is defined, we mean synonymously that there exists a constant

$C \geq 0$ such that $|f(x)| \leq Cg(x)$ for all $x \in X$. The ‘implied constant’ is any admissible value of C . It may depend on the set X which is always specified or clear in context. The notation $f \asymp g$ means $f \ll g$ and $g \ll f$. On the other hand $f(x) = o(g(x))$ as $x \rightarrow x_0$ is a topological statement meaning that $f(x)/g(x) \rightarrow 0$ as $x \rightarrow x_0$. We also use the $O()$ notation in other types of expressions; the meaning should be clear: e.g., $f(x) \leq g(x) + O(h(x))$ for $x \in X$, means that $f \leq g + h_1$ in X for some (non-negative) function h_1 such that $h_1 = O(h)$. (For instance, $x \leq x^2 + O(1)$ for $x \geq 1$, but it is not true that $x - x^2 = O(1)$.)

In this book, any statement of a lemma, proposition, theorem or corollary will include an explicit mention of which parameters the ‘implied constant’ depends on; any divergence from this principle is an error, and the author should be made aware of it. The same explicitness will be true for many, but not all, of the intermediate statements (where sometimes it will be clear enough what the parameters involved are, from the flow of the argument). This insistence may look pedantic, but uniformity in parameters is crucial to many applications of analytic number theory, and this should make the text usable by all mathematicians with confidence that there is no hidden dependency. (Algebraic-minded readers may note that indicating the dependency of those parameters is somewhat analogous to stating explicitly in which category a morphism between two objects is defined; the author’s experience is that not having this information clearly stated *even if it is completely obvious for knowledgeable readers* can create a lot of confusion for beginners.)

For a group G , G^\sharp denotes the set of its conjugacy classes, and for a conjugacy-invariant subset $X \subset G$, $X^\sharp \subset G^\sharp$ is the corresponding set of conjugacy classes. The conjugacy class of $g \in G$ is denoted g^\sharp .

For q a power of a prime number, \mathbf{F}_q denotes a finite field with q elements.

Unless otherwise specified (as in Chapter 5), p always denotes a prime number. If $n \geq 1$ is an integer, sums or products over divisors of n always mean divisors $d \geq 1$. We use standard arithmetic functions φ , ψ , ω and μ ,³ defined as follows for an integer $n \geq 1$ in terms of the prime factors of n :

$$\varphi(n) = n \prod_{p|n} (1 - p^{-1}), \quad \psi(n) = n \prod_{p|n} (1 + p^{-1}), \quad \omega(n) = |\{p \mid p \mid n\}|,$$

$$\mu(n) = \begin{cases} (-1)^k & \text{if } n = p_1 \cdots p_k \text{ with } p_1 < \cdots < p_k, \\ 0 & \text{otherwise,} \end{cases}$$

³ No confusion should arise with measures also denoted μ .

We denote as (a, b) the greatest common divisor of integers a and b , unless this creates ambiguity with pairs of integers. Similarly, $[a, b]$ is the least common multiple. An integer $n \geq 1$ is *squarefree* if it is not divisible by the square of a prime p , or equivalently if $\mu(n) \neq 0$. We use the shorthand notation

$$\sum_m^b \alpha(m)$$

for a sum restricted to squarefree integers m .

We denote by $\pi(x)$ the prime counting function, i.e., the number of primes $p \leq x$, and by $\pi(x; q, a)$ the prime counting function in arithmetic progressions, i.e., the number of primes $p \leq x$ which are congruent to a modulo q . Of course, $\pi(x; q, a)$ is bounded if and only if $(a, q) \neq 1$ (by Dirichlet's theorem on primes in arithmetic progressions).

We recall some asymptotic formulas of prime number theory, the second of which is a strong form of the Prime Number Theorem:

$$\sum_{p \leq x} \frac{1}{p} = \log \log x + O(1), \quad \pi(x) = \frac{x}{\log x} + O\left(\frac{x}{(\log x)^2}\right),$$

for $x \geq 3$.

For $z \in \mathbf{C}$, we denote $e(z) = \exp(2i\pi z)$, so that $e(\cdot)$ is a non-trivial homomorphism $\mathbf{C}/\mathbf{Z} \rightarrow \mathbf{C}^\times$.

In probabilistic contexts, $\mathbf{P}(A)$ is the probability of an event, $\mathbf{E}(X)$ is the expectation of a random variable X , $\mathbf{V}(X)$ its variance, and $\mathbf{1}_A$ is the characteristic function of an event A . See Appendix F for the basic definitions.

Let k be a field, and V a k -vector space of even dimension $\dim V = 2g$. If $\langle \cdot, \cdot \rangle : V \times V \rightarrow k$ is a non-degenerate alternating bilinear form on V , we denote by $Sp(V)$, $Sp(\langle \cdot, \cdot \rangle)$ or more commonly by $Sp(2g, k)$ the *symplectic group* of V , namely the group of invertible linear transformations of V preserving this bilinear form; it is the group of those $g \in GL(V)$ such that

$$\langle gv, gw \rangle = \langle v, w \rangle$$

for all $v, w \in V$. The notation $Sp(2g, k)$ is justified by the fact that, up to isomorphism, there is only one non-degenerate alternating bilinear form on V . If a specific model is needed, one can fix a vector space W of dimension g , and put $V = W \oplus W'$, where W' is the dual of W , and let

$$\langle (v_1, \ell_1), (v_2, \ell_2) \rangle = \ell_1(v_2) - \ell_2(v_1).$$

The subspaces W and W' are then instances of *Lagrangian subspaces*, i.e., subspaces of maximal dimension g such that the restriction of the alternating form to the subspace is identically zero. All Lagrangian subspaces of V are

images of any fixed one (such as W above) by an element of $Sp(V)$, i.e., $Sp(V)$ acts transitively on the set of Lagrangian subspaces. If W_1, W_2 are Lagrangian subspaces, they are *transverse* if $W_1 \cap W_2 = 0$, or equivalently if both together span V .

Moreover, we denote by $CSp(V)$, $CSp(\langle \cdot, \cdot \rangle)$ or $CSp(2g, k)$ the group of *symplectic similitudes*, i.e., of those $g \in GL(V)$ such that

$$\langle gv, gw \rangle = m(g)\langle v, w \rangle$$

for all $v, w \in V$, where $m(g) \in k^\times$ is a scalar called the *multiplicator* of g . This is a surjective group homomorphism, and there is therefore an exact sequence

$$1 \rightarrow Sp(V) \rightarrow CSp(V) \xrightarrow{m} k^\times \rightarrow 1.$$

We recall the formulas for the cardinality of $GL(n, \mathbf{F}_q)$ and $Sp(2g, \mathbf{F}_q)$ for a finite field \mathbf{F}_q with q elements:

$$|GL(n, \mathbf{F}_q)| = \prod_{k=0}^{n-1} (q^n - q^k) = q^{n(n-1)/2} \prod_{k=1}^n (q^k - 1), \quad (0.1)$$

$$|Sp(2g, \mathbf{F}_q)| = q^{g^2} \prod_{k=1}^g (q^{2k} - 1). \quad (0.2)$$

When working with matrices $g \in M(n, A)$, where A is a commutative ring with unit, we will consider both the standard *characteristic polynomial* of g , namely $\det(T - g) \in A[T]$, which is a monic polynomial of degree n taking value $(-1)^n \det(g)$ at 0; and the *reversed characteristic polynomial* $\det(\text{Id} - Tg) \in A[T]$, where Id is the identity matrix. This is of degree equal to the rank of g , takes value 1 at 0, and has leading term $\det(g)T^n$ if g is invertible. Obviously, whenever invertible matrices are considered, all results on either of these can be restated in terms of the other, or of $\det(g - T)$: we have

$$\det(\text{Id} - gT) = T^n \det(T^{-1} - g).$$

If we wish to speak of the characteristic polynomial of an endomorphism of a free A -module V of finite rank, we write $\det(T - A \mid V)$ or $\det(\text{Id} - TA \mid V)$.

If G is a group, $[G, G]$ is the commutator subgroup, generated by commutators $[x, y] = xyx^{-1}y^{-1}$ for $x, y \in G$, and the abelian group $G/[G, G]$ is the *abelianization* of G .

The symmetric group on n letters is denoted \mathfrak{S}_n . Moreover, for $g \geq 1$, W_{2g} denotes the group of *signed permutations* of g pairs $(2i - 1, 2i)$, $1 \leq i \leq 2g$,

i.e., the subgroup of elements $\sigma \in \mathfrak{S}_{2g}$ such that $\sigma(\{2i-1, 2i\})$ is a pair $\{2j, 2j-1\}$ for all i . This group has order $2^g g!$ and sits in an exact sequence

$$1 \rightarrow \{\pm 1\}^g \rightarrow W_{2g} \xrightarrow{p} \mathfrak{S}_g \rightarrow 1,$$

where the right-hand map assigns to $\sigma \in W_{2g}$ the permutation of the g pairs $(2i-1, 2i)$, the natural generators σ_i of the kernel being the signed permutations which act as the identity except for $\sigma(2i-1) = 2i$, $\sigma(2i) = 2i-1$.

1

Introduction

1.1 Presentation

Classical sieve theory is concerned with the problem of the asymptotic evaluation of averages of arithmetic functions over integers constrained by congruence restrictions modulo a set of primes. Often the function in question is the characteristic function of some interesting sequence and the congruence restrictions are chosen so that those integers remaining after the sieving process are, for instance, primes or ‘almost’ primes.

If the congruence conditions are phrased as stating that the only integers n which are allowed are those with reduction modulo a prime p not in a certain set Ω_p , then a familiar dichotomy arises: if Ω_p contains few residue classes (typically, a bounded number as p increases), the setting is that of a ‘small’ sieve. The simplest such case is the detection of primes with $\Omega_p = \{0\}$. If, on the other hand, the size of Ω_p increases with p , the situation is that of a ‘large’ sieve. The first such sieve was devised by Linnik to investigate the question of Vinogradov of the size of the smallest quadratic non-residue modulo a prime.

There have already been a number of works extending ‘small’ sieves to more general situations, where the objects being sifted are not necessarily integers. One may quote among these the vector sieve of Brüdern and Fouvry [17], with applications to Lagrange’s theorem with almost prime variables; the ‘crible étrange’ of Fouvry and Michel [42], with applications to sign changes of Kloosterman sums, and Poonen’s striking sieve procedure for finding smooth hypersurfaces of large degree over finite fields [105] (which we describe briefly in Example 4.11).

Similarly, the large sieve has been extended in some ways, in particular (quite early on) to deal with sieves in \mathbf{Z}^d , $d \geq 1$, or in number fields (see, e.g. [46]). Interesting applications have been found, e.g. Duke’s theorem on elliptic curves over \mathbf{Q} with ‘maximal’ p -torsion fields for all p [32]. All these were much of

the same flavour however, and in particular depended only on the character theory of finite abelian groups as far as the underlying harmonic analysis was concerned.

In [80], we introduced a new large sieve inequality to study the average distribution of Frobenius conjugacy classes in the monodromy groups of a family (\mathcal{F}_ℓ) of \mathbf{F}_ℓ -adic sheaves on a variety over a finite field. Although the spirit of the large sieve is clearly recognizable, the setting is very different, and the harmonic analysis involves both non-abelian finite groups and the deep results of Deligne on the Riemann Hypothesis over finite fields. Our first application of this new sieve was related to the ‘generic’ arithmetic behaviour of the numerator of the zeta function of a smooth projective curve in a family with large monodromy, improving significantly a result of Chavdarov [22]. (We will survey and again improve these results in Chapter 8.)

As explained in the preface, while working on devising a general framework of the sieve that can recover both the classical forms or the version in [80], a number of new applications emerged. Some of them are in areas of number theory not usually directly linked to sieve methods, and some in decidedly different contexts. Hence the goal of this book is to present the large sieve as a general mathematical *principle* which has potential applications outside number theory. For this reason, we start from scratch, assuming only a knowledge of basic linear algebra and properties of finite-dimensional Hilbert spaces to derive the basic inequality.

Roughly speaking, this inequality states that, given a measure space X with finite measure, and surjective maps from X to a family (X_ℓ) of *finite* sets, the measure of the set of those $x \in X$ which have image in X_ℓ outside some given sets Ω_ℓ , for finitely many ℓ , can be estimated from above by means of two quantities. One involves the ‘densities’ of the sets Ω_ℓ in X_ℓ , and is independent of X , while the other (the ‘large sieve constant’) is the norm of a certain bilinear form which depends on X and X_ℓ , but is independent of Ω_ℓ . This form of the sieve statement is similar to Montgomery’s inequality, and much stronger than Linnik’s original version (see, e.g. [98], [11], [67, 7.4]).

Obtaining this inequality is really straightforward and is done, in Chapter 2, in a few pages – the innovation, for what it’s worth, is in working in the generality we consider. This does not by itself prove anything, because the large sieve constant needs to be estimated before applications can be derived, and the estimation may turn out to be impossible, or trivial. However, the problem turns out to be further reducible to the study of certain ‘exponential sums’ (or integrals) over X , which suggests that strong estimates should exist in many situations, related to the equidistribution of the image of X in X_ℓ . This equidistribution may be expected to be true in many cases, for fixed ℓ at least, but a key issue is

uniformity with respect to ℓ : an explicit form of the error term in the equidistribution is required to proceed. In the classical case, the bilinear form estimate was first considered by Bombieri and given its most general expression by Davenport and Halberstam.

This is the time to discuss a thorny terminological issue: this inequality (in its most refined version) takes the form

$$\sum_r \left| \sum_{M \leq n < M+N} a_n e(n\xi_r) \right|^2 \leq (N-1+\delta^{-1}) \sum_n |a_n|^2 \quad (1.1)$$

for arbitrary complex numbers a_n and ‘angles’ $\xi_r \in \mathbf{R}/\mathbf{Z}$ which are δ -spaced (i.e., such that $\min_{n \in \mathbf{Z}} |\xi_r - \xi_s - n| \geq \delta$ for $r \neq s$). It is often itself called ‘the large sieve inequality’, although it does not mention any idea of sieve, because of its link with the proof of Montgomery’s inequality. Correspondingly, when generalizations of (1.1) were developed for independent reasons (replacing the characters $x \mapsto e(x\xi_r)$ by other functions), they were also called ‘large sieve inequalities’, even when any link to sieve theory had utterly vanished. And in fact these inequalities, particularly those involving Fourier coefficients of automorphic forms of various types, form an important body of work which has had tremendous applications in analytic number theory, starting with the work of Iwaniec, and Deshouillers–Iwaniec, and later with variants due to Duke, Duke–Kowalski, Venkatesh and others. We will not say anything beyond this, and we refer to [67, Section 7.7] for a short survey with some applications.

After presenting and commenting on the basic framework, the rest of the book is devoted to the explanation of a number of instances of sieves and the issues surrounding them. This is done first with the examples of Chapter 4 which present a number of (mostly) classical situations in this context, and describe some of their applications for convenience. We also indicate there the relation with the inclusion-exclusion technique in probability and combinatorics, which shows in particular that the general sieve bound is sharp, and include a first new application: an amusing ‘elliptic sieve’ which is related to questions surrounding the number of prime divisors of the denominators of rational points on an elliptic curve. In turn, this is linked to the analysis of the prime factorization of elements of the so-called ‘elliptic divisibility sequences’ first introduced by M. Ward. We find rather easily that ‘most’ elements have many prime factors, which complements recent heuristics and results of Silverman, Everest, Ward and others concerning the paucity of primes and prime powers in such sequences.

The following chapters are less classical and concern new (or recent) applications of the sieve ideas, which are quite independent of one another.

‘Probabilistic’ sieves are discussed briefly in Chapter 6, with an application to ‘random’ finitely presented groups, and sieving in a discrete finitely generated group G is described in much more detail in Chapter 7, where some of the most appealing new results are obtained. Indeed, for symmetric random walks on some finitely generated groups, a very transparent treatment of the large sieve constant is possible, and Property (τ) (or the expanding properties of Cayley graphs of quotients of G) appears as a completely natural tool. When this feature is present, it leads to strong sieve results. Moreover, very interesting applications arise, including surprising ones in geometry or topology.

Finally, in Chapter 8, we review and extend the sieve result of [80] concerning the distribution of geometric Frobenius conjugacy classes in finite monodromy groups over finite fields, and derive some new applications. There are links here with the case of arithmetic groups, and comparison of the sieve bounds coming from Property (τ) in the former case and the Riemann Hypothesis over finite fields in the latter is quite interesting.

The final part of the book is a series of appendices which review briefly some of the topics which are probably not known to all readers. This includes a discussion of small sieves, for purpose of comparison and reference, including a sample application; a survey of some techniques that are used to prove density results in matrix groups over finite fields, which are also of independent interest and involve work of Chavdarov [22] and non-trivial estimates for exponential sums over finite fields; a survey of representation theory of groups, involving both the classical theory for finite groups, and what is needed to describe Property (T) and Property (τ) ; some estimates for sums of multiplicative functions; and a short survey of basic topological facts which we use in some of our applications.

Whenever we treat an example, we give at least all definitions required to understand the essential parts of the statements, and precise references for any unproved facts which can not be assumed to be known by every potential reader. It is expected that most readers will at least once think ‘Everyone knows *this*!’ when reading some part of the notes, but they may not be able to say this of all such basic references.

1.2 Some new applications of the large sieve

Before going further, it seems natural to list here a few applications of the sieve framework we are going to describe. Most of those below are, to the best of our knowledge, new results, although some of them could well have been

proven before. We seek concreteness in this list: the precise results will usually be stronger and more general.

Our first result is in fact obtained from the ‘traditional’ large sieve in one variable, which we apply in a rather twisted way.

Theorem 1.1 *Let E/\mathbf{Q} be an elliptic curve with rank $r \geq 1$ given by a Weierstrass equation*

$$y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6, \quad \text{where } a_i \in \mathbf{Z}.$$

For $x \in E(\mathbf{Q})$, let $\omega_E(x)$ be the number of primes, without multiplicity, dividing the denominator of the coordinates of x , with $\omega_E(0) = +\infty$. Let $h(x)$ denote the canonical height on E .

Then for any fixed real number κ with $0 < \kappa < 1$, we have

$$\frac{|\{x \in E(\mathbf{Q}) \mid h(x) \leq T \text{ and } \omega_E(x) < \kappa \log \log T\}|}{|\{x \in E(\mathbf{Q}) \mid h(x) \leq T\}|} \ll (\log \log T)^{-1},$$

for $T \geq 3$, where the implied constant depends only on E and κ .

The second statement is an example of the philosophy that random walks on a set give a way of stating properties of random elements of X , even when there is no natural probability measure on X . Here X is the set of integers \mathbf{Z} , and we use simple random walks to compensate for the absence of a translation-invariant probability measure on \mathbf{Z} .

Theorem 1.2 *Let (S_n) be a simple random walk on \mathbf{Z} , i.e.,*

$$S_n = X_1 + \cdots + X_n$$

where (X_k) is a sequence of independent random variables with $\mathbf{P}(X_k = \pm 1) = 1/2$ for all k .

Let $\varepsilon > 0$ be given, $\varepsilon \leq 1/4$. For any odd $q \geq 1$, any a coprime with q , we have

$$\mathbf{P}(S_n \text{ is prime and } \equiv a \pmod{q}) \ll \frac{1}{\varphi(q)} \frac{1}{\log n}$$

if $n \geq 1$, $q \leq n^{1/4-\varepsilon}$, the implied constant depending only on ε .

This is proved in Chapter 6. It may be expected that results of this type can be recovered from their ‘deterministic’ analogues using the Central Limit Theorem. However, this is not likely to be feasible (or wise) when considering similar questions about random unimodular matrices. In Chapter 7, we prove the following result using Property (τ) , which confirms that generic elements of $SL(n, \mathbf{Z})$ have ‘arithmetically generic’ characteristic polynomials:

Theorem 1.3 *Let $n \geq 2$ be an integer, let $G = SL(n, \mathbf{Z})$ and let $S = S^{-1} \subset G$ be a finite generating set of G , e.g., the finite set of elementary matrices with ± 1 entries off the diagonal. Let (X_k) be the simple left-invariant random walk on G , i.e., a sequence of G -valued random variables such that $X_0 = 1$ and*

$$X_{k+1} = X_k \xi_{k+1} \text{ for } k \geq 0,$$

where (ξ_k) is a sequence of S -valued independent random variables with

$$\mathbf{P}(\xi_k = s) = \frac{1}{|S|} \quad \text{for all } s \in S.$$

Then, almost surely, there are only finitely many k for which the characteristic polynomial $\det(T - X_k) \in \mathbf{Z}[T]$ does not have the full symmetric group \mathfrak{S}_n as Galois group, or in other words, the set of matrices in $SL(n, \mathbf{Z})$ with characteristic polynomials having small Galois group is transient for the random walk. In particular, so is the set of those having reducible characteristic polynomial.

In fact (see Theorem 7.4), we will derive this by showing that the probability that $\det(T - X_k)$ be reducible decays exponentially fast with k (in the case $n \geq 3$ at least). The following is a consequence of a similar statement for symplectic groups, and it answers a question of Maher [96, Question 1.3] (see Proposition 7.17).

Theorem 1.4 *Let $g \geq 1$ be an integer, let G be the mapping class group of a closed surface Σ_g of genus g . Then the set of non-pseudo-Anosov elements in G is transient for any symmetric random walk on G where the steps are chosen among a fixed finite symmetric generating set of G .*

These two examples of sieves in discrete groups correspond to properties which are invariant under conjugation. The next result does not have this property, showing that the sieve is not limited to this situation. For the sake of diversity, we state the result somewhat differently in the language of products of N matrices chosen among the generating set.

Theorem 1.5 *Let $n \geq 3$ be an integer, let $G = SL(n, \mathbf{Z})$, and let $S = S^{-1} \subset G$ be a finite symmetric generating set. Then there exists $\beta > 0$ such that for any $N \geq 1$, we have*

$$|\{w \in S^N \mid \text{one entry of the matrix } g_w \text{ is a square}\}| \ll |S|^{N(1-\beta)},$$

where $g_w = s_1 \cdots s_N$ for $w = (s_1, \dots, s_N) \in S^N$, and β and the implied constant depend only on n and S .

Finally, here is a sample of what the sieve for Frobenius can do, as described in Chapter 8. Except for a slightly weaker exponent γ , it could have been proved easily with the techniques of [80].

Theorem 1.6 *Let q be a power of a prime number $p \geq 5$, $g \geq 1$ an integer and let $f \in \mathbf{F}_q[T]$ be a squarefree polynomial of degree $2g$. For t not a zero of f , let C_t denote the smooth projective model of the hyperelliptic curve*

$$y^2 = f(x)(x - t),$$

and let J_t denote its Jacobian variety. Then we have

$$|\{t \in \mathbf{F}_q \mid f(t) \neq 0 \text{ and } |C_t(\mathbf{F}_q)| \text{ is a square}\}| \ll gq^{1-\gamma}(\log q),$$

$$|\{t \in \mathbf{F}_q \mid f(t) \neq 0 \text{ and } |J_t(\mathbf{F}_q)| \text{ is a square}\}| \ll gq^{1-\gamma}(\log q)$$

where $\gamma = (4g^2 + 2g + 4)^{-1}$, and the implied constants are absolute.

2

The principle of the large sieve

2.1 Notation and terminology

We will start by describing a very general type of sieve. The goal is to reach an analogue of the large sieve inequality, in the sense of a reduction of a sieve bound to a bilinear form estimate.

We start by introducing the notation and terminology. Many readers, especially analytic number theorists, may find it excessively formal, but the framework we describe has so many different incarnations that it seems preferable to be very precise in this book, and to give a name to the objects involved to refer to them later on. Concrete applications will be able to eschew reproducing all this, by using self-contained statements such as those included in Section 3.5, which involve none of the newfangled terminology.

Hence, let's start. First of all, the *sieve setting* is a triple $\Psi = (Y, \Lambda, (\rho_\ell))$ consisting of

- a set Y ;
- an index set Λ ;
- for all $\ell \in \Lambda$, a surjective map $\rho_\ell : Y \rightarrow Y_\ell$ where Y_ℓ is a finite set.

In combinatorial terms, this might be thought of as a family of colourings of the set Y . In applications, Λ will often be a subset of primes (or prime ideals in some number field), but as first pointed out by Zywinia, this is not necessary for the formal part of setting up the sieve, and although the generality is not really abstractly greater, it is convenient to allow arbitrary Λ .

Then, a *siftable set* associated to $\Psi = (Y, \Lambda, (\rho_\ell))$ is a triple $\Upsilon = (X, \mu, F)$ consisting of

- a measure space (X, μ) with $\mu(X) < +\infty$;
- a map $F : X \rightarrow Y$ such that the composites $X \rightarrow Y \rightarrow Y_\ell$ are measurable, i.e., the sets $\{x \in X \mid \rho_\ell(F_x) = y\}$ are measurable for all ℓ and all $y \in Y_\ell$.

The simplest case is when X is a finite set and μ is counting measure. We call this the *counting case*. Even when this is not the case, for notational convenience, we will usually write $|B|$ for the measure $\mu(B)$ of a measurable set $B \subset X$.

The last pieces of data are a finite subset \mathcal{L}^* of Λ , called the *prime sieve support*, and a family $\Omega = (\Omega_\ell)$ of *sieving sets*,¹ $\Omega_\ell \subset Y_\ell$, defined for $\ell \in \mathcal{L}^*$.

With this final data $(\Psi, \Upsilon, \mathcal{L}^*, \Omega)$, we can define the sieve problem.

Definition 2.1 Let $\Psi = (Y, \Lambda, (\rho_\ell))$ be a sieve setting, $\Upsilon = (X, \mu, F)$ a siftable set, \mathcal{L}^* a prime sieve support and Ω a family of sieving sets. Then the sifted sets are

$$\begin{aligned} S(Y, \Omega; \mathcal{L}^*) &= \{y \in Y \mid \rho_\ell(y) \notin \Omega_\ell \text{ for all } \ell \in \mathcal{L}^*\}, \\ S(X, \Omega; \mathcal{L}^*) &= \{x \in X \mid \rho_\ell(F_x) \notin \Omega_\ell \text{ for all } \ell \in \mathcal{L}^*\}. \end{aligned}$$

The latter is also $F^{-1}(S(Y, \Omega; \mathcal{L}^*))$ and is a measurable subset of X .

The problem we will consider is to find estimates for the measure $|S(X, \Omega; \mathcal{L}^*)|$ of the sifted set. Here we have in mind that the sieve setting is fixed, while there usually will be an infinite sequence of siftable sets with size $|X|$ going to infinity; this size will be the main variable in the estimates.

Example 2.2 The classical sieve arises as follows: the sieve setting is

$$\Psi = (\mathbf{Z}, \{\text{primes}\}, \mathbf{Z} \rightarrow \mathbf{Z}/\ell\mathbf{Z})$$

and the siftable sets are $X = \{n \mid M < n \leq M+N\}$ with counting measure and $F_x = x$ for $x \in X$. Then the sifted sets become the classical sets of integers in an interval with reductions modulo primes in \mathcal{L}^* lying outside a subset $\Omega_\ell \subset \mathbf{Z}/\ell\mathbf{Z}$ of residue classes.

In most cases, (X, μ) will be a finite set with counting measure, and often $X \subset Y$ with $F_x = x$ for $x \in X$. See Chapter 8 for a conspicuous example where F is not the identity, Chapter 6 for interesting situations where the measure space (X, μ) is a probability space, and F a random variable, and Chapter 7 for another example.

2.2 The large sieve inequality

We will now indicate one type of inequality that reduces the sieve problem to the estimation of a *large sieve constant* Δ . The latter is a more analytic problem,

¹ Sometimes, Ω will also denote a probability space, but no confusion should arise.

and can be attacked in a number of ways. This large sieve constant depends on most of the data involved, but is independent of the sieving sets.

First we need some more notation. Given a sieve setting Ψ , we let $S(\Lambda)$ denote the set of finite subsets $m \subset \Lambda$. In order to simplify notation, since $S(\Lambda)$ may be identified with the set of squarefree integers $m \geq 1$ in the classical case where Λ is the set of primes, we write $\ell \mid m$ for $\ell \in m$ when $\ell \in \Lambda$ and $m \in S(\Lambda)$ (and similarly for $n \mid m$ instead of $n \subset m$ if $n, m \in S(\Lambda)$). Also we sometimes do not explicitly distinguish between $\ell \in \Lambda$ and $\{\ell\} \in S(\Lambda)$.

A sieve support \mathcal{L} associated to a prime sieve support \mathcal{L}^* is any (finite) family of subsets of \mathcal{L} . (In general, \mathcal{L} will have additional properties, in particular it will be such that $\{\ell\} \in \mathcal{L}$ for any $\ell \in \mathcal{L}^*$, but it is not necessary to assume this.)

If Λ is a set of primes, \mathcal{L} ‘is’ a set of squarefree integers only divisible by primes in \mathcal{L}^* (including possibly $m = 1$, not divisible by any prime).

For $m \in S(\Lambda)$, let

$$Y_m = \prod_{\ell \mid m} Y_\ell$$

and let $\rho_m : Y \rightarrow Y_m$ be the obvious product map. (In other words, we look at all ‘refined’ colourings of Y obtained by looking at all possible finite tuples of colourings.) If $m = \emptyset$, Y_m is a set with a single element, and ρ_m is a constant map. Note that ρ_m is not surjective in general.

We will consider functions on the various sets Y_m , and it will be important to endow the space of complex-valued functions on Y_m with appropriate and consistent inner products. For this purpose, we assume given for $\ell \in \Lambda$ a density

$$\nu_\ell : Y_\ell \rightarrow [0, 1]$$

(often denoted simply by ν when no ambiguity is possible) such that the inner product on functions $f : Y_\ell \rightarrow \mathbf{C}$ is given by

$$\langle f, g \rangle = \sum_{y \in Y_\ell} \nu_\ell(y) f(y) \overline{g(y)}.$$

We assume that $\nu(y) > 0$ for all $y \in Y_\ell$, in order that this hermitian form be positive definite (it will be clear that $\nu(y) \geq 0$ would suffice, with minor changes, but the stronger assumption is no problem for applications), and that ν is a probability density, i.e., we have

$$\sum_{y \in Y_\ell} \nu_\ell(y) = 1. \tag{2.1}$$

We denote by $L^2(Y_\ell, \nu_\ell)$, or simply $L^2(Y_\ell)$, the Hilbert space of functions on Y_ℓ with this inner product.

Now, using the product structure, we define densities ν_m and corresponding inner products on the spaces of functions $Y_m \rightarrow \mathbf{C}$: we have

$$\nu_m(y) = \prod_{\ell|m} \nu_\ell(y_\ell)$$

for $y = (y_\ell) \in Y_m$, and

$$\langle f, g \rangle = \sum_{y \in Y_m} \nu_m(y) f(y) \overline{g(y)}$$

for $f, g : Y_m \rightarrow \mathbf{C}$. In particular, Property (2.1) still holds, and if f, g are functions of the type

$$f = \otimes_{\ell|m} f_\ell, \quad g = \otimes_{\ell|m} g_\ell$$

(which means, e.g., that $f(y) = \prod f_\ell(y_\ell)$ for $y = (y_\ell) \in Y_m$), we have

$$\langle f, g \rangle = \prod_{\ell|m} \langle f_\ell, g_\ell \rangle,$$

with the inner products on Y_ℓ on the right-hand side.

We will interpret ν_ℓ or ν_m as measures on Y_ℓ or Y_m (so that ν_m is then the product measure of the ν_ℓ for $\ell \mid m$), and often drop the subscript, so we will write for instance

$$\nu(\Omega_\ell) = \sum_{y \in \Omega_\ell} \nu(y), \quad \text{for } \Omega_\ell \subset Y_\ell.$$

We denote by $L^2(Y_m, \nu_m)$ or $L^2(Y_m)$ the space of complex-valued functions on Y_m with the inner product thus defined.

The simplest example (uniform density) is when $\nu_\ell(y) = 1/|Y_\ell|$, so that $\nu(y) = 1/|Y_m|$ for all m , but we will see, e.g. in Chapter 3, important natural cases where ν is not uniform. It will be clear in the remarks and chapters following the statement of the sieve inequality that, in general, the apparent freedom of choice of ν_ℓ is illusory (only one choice will lead to good results for a given sieve setting).

Having chosen the density, assume given, for any $\ell \in \Lambda$, an orthonormal basis \mathcal{B}_ℓ of $L^2(Y_\ell, \nu_\ell)$, such that the constant function 1 is in \mathcal{B}_ℓ , and let $\mathcal{B}_\ell^* = \mathcal{B}_\ell - \{1\}$. (Equivalently, \mathcal{B}_ℓ^* is an orthonormal basis of the ‘primitive’ subspace

$$L_0^2(Y_\ell) = \left\{ f \in L^2(Y_\ell) \mid \langle f, 1 \rangle = \sum_y \nu(y) f(y) = 0 \right\}$$

which is the orthogonal of the constant functions on Y_ℓ .) For $m \in S(\Lambda)$, define

$$\mathcal{B}_m = \prod_{\ell|m} \mathcal{B}_\ell, \quad \mathcal{B}_m^* = \prod_{\ell|m} \mathcal{B}_\ell^*,$$

and identify elements of \mathcal{B}_m and \mathcal{B}_m^* with functions on Y_m , the function corresponding to $(\varphi_\ell) \in \mathcal{B}_m$ being given by

$$(y_\ell) \mapsto \prod_{\ell|m} \varphi_\ell(y_\ell);$$

for $m = \emptyset$, the consistent definition is to let $\mathcal{B}_m = \mathcal{B}_m^* = \{1\}$.

Note that \mathcal{B}_m is an orthonormal basis of $L^2(Y_m, \nu_m)$, and \mathcal{B}_m^* is an orthonormal basis of a certain ‘primitive’ subspace, identified naturally with the tensor product

$$L_0^2(Y_m) = \bigotimes_{\ell|m} L_0^2(Y_\ell). \quad (2.2)$$

Here now is the first sieve inequality, which in the classical case was first formulated by Montgomery.

Proposition 2.3 *Let $\Psi, \Upsilon, \mathcal{L}^*$ be as above. For any sieve support \mathcal{L} associated to \mathcal{L}^* , i.e., any finite family of subsets of \mathcal{L} , let $\Delta = \Delta(X, \mathcal{L})$ denote the large sieve constant, which is by definition the smallest non-negative real number such that*

$$\sum_{m \in \mathcal{L}} \sum_{\varphi \in \mathcal{B}_m^*} \left| \int_X \alpha(x) \varphi(\rho_m(F_x)) d\mu(x) \right|^2 \leq \Delta \int_X |\alpha(x)|^2 d\mu(x) \quad (2.3)$$

for any square integrable function $\alpha : X \rightarrow \mathbf{C}$.

Then for arbitrary sieving sets $\Omega = (\Omega_\ell)$, we have

$$|S(X, \Omega; \mathcal{L}^*)| \leq \Delta H^{-1}$$

where

$$H = \sum_{m \in \mathcal{L}} \prod_{\ell|m} \frac{\nu(\Omega_\ell)}{\nu(Y_\ell - \Omega_\ell)} = \sum_{m \in \mathcal{L}} \prod_{\ell|m} \frac{\nu(\Omega_\ell)}{1 - \nu(\Omega_\ell)}. \quad (2.4)$$

Example 2.4 In the classical case, with $Y = \mathbf{Z}$ and $Y_\ell = \mathbf{Z}/\ell\mathbf{Z}$, we can identify Y_m with $\mathbf{Z}/m\mathbf{Z}$ by the Chinese Remainder Theorem. With $\nu(y) = 1/\ell$ for all ℓ and all y , we have $\nu_m(y) = 1/m$ for all squarefree m . The orthonormal basis of functions on Y_ℓ (for the uniform density) used in the classical large sieve is provided by the additive characters

$$\begin{cases} \mathbf{Z}/\ell\mathbf{Z} & \longrightarrow \mathbf{C}^\times \\ x & \mapsto e\left(\frac{ax}{\ell}\right) \end{cases}$$

where a runs also over $\mathbf{Z}/\ell\mathbf{Z}$. It is then easy to check that the corresponding orthonormal basis \mathcal{B}_m of $\mathbf{Z}/m\mathbf{Z}$ can be identified with that of additive characters

$$x \mapsto e\left(\frac{ax}{m}\right)$$

where now $a \in \mathbf{Z}/m\mathbf{Z}$. This is really part of the general theory (see Section 3.1), but can be checked directly from scratch. Indeed, the analytic expression for the Chinese Remainder Theorem in $\mathbf{Z}/m\mathbf{Z}$ is

$$z \equiv \sum_{\ell|m} (m/\ell) \overline{(m/\ell)} z_\ell \pmod{m}$$

for $z \in \mathbf{Z}/m\mathbf{Z}$ such that $z \equiv z_\ell$ for $\ell \mid m$, where $\overline{(m/\ell)}$ is the modular inverse modulo ℓ (it is clear that the right-hand side is well-defined modulo m , and is indeed congruent to z_ℓ modulo ℓ for $\ell \mid m$). Thus, denoting by a_ℓ and x_ℓ the components of a and x modulo ℓ , we have

$$e\left(\frac{ax}{m}\right) = e\left(\sum_{\ell|m} \frac{\overline{(m/\ell)} a_\ell x_\ell}{\ell}\right) = \prod_{\ell|m} e\left(\frac{\overline{(m/\ell)} a_\ell x_\ell}{\ell}\right),$$

showing that the additive characters in question are in \mathcal{B}_m . Since their number is m also, the claim holds.

It is also easy to check that such a character belongs to \mathcal{B}_m^* if and only if a and m are coprime. Indeed, in the representation above, a does *not* correspond to an element in \mathcal{B}_m^* if and only if $a_\ell \overline{(m/\ell)} \equiv 0 \pmod{\ell}$ for some $\ell \mid m$, which is equivalent with $a_\ell \equiv 0 \pmod{\ell}$, hence to $\ell \mid a$ for some $\ell \mid m$.

Remark 2.5 The large sieve constant as defined in Proposition 2.3 is independent of the choices of bases \mathcal{B}_ℓ (containing the constant function 1). Here is a more intrinsic definition which shows this, and provides a first hint of the link with classical (small) sieve axioms. It's not clear how useful this intrinsic definition can be in practice, which explains why we kept a concrete version in the statement of Proposition 2.3.

By definition, the inequality (2.3) means that Δ is the square of the norm of the linear operator

$$T \begin{cases} L^2(X, \mu) & \longrightarrow & \bigoplus_{m \in \mathcal{L}} L_0^2(Y_m)' \\ \alpha & \longmapsto & \left(f \mapsto \int_X \alpha(x) f(\rho_m(F_x)) d\mu(x) \right)_m \end{cases}$$

where the direct sum over m is orthogonal and $L_0^2(Y_m)'$ is the space of linear functionals on the primitive subspace (2.2), with the usual norm

$$\|f^*\| = \max_{f \neq 0} \frac{|\langle f^*, f \rangle|}{\|f\|}.$$

Since we are dealing with Hilbert spaces, $L_0^2(Y_m)'$ is canonically isometric to $L_0^2(Y_m)$, and Δ is the square of the norm of the operator

$$T_1 \begin{cases} L^2(X, \mu) & \longrightarrow \bigoplus_{m \in \mathcal{L}} L_0^2(Y_m) \\ \alpha & \longmapsto T_1(\alpha) \end{cases}$$

where $T_1(\alpha)$ is the vector such that $\langle f, T_1(\alpha) \rangle = T(\alpha)(f)$ for $f \in L_0^2(Y_m)$, $m \in \mathcal{L}$. This vector is easy to identify: we have

$$\int_X \alpha(x) f(\rho_m(F_x)) d\mu(x) = \sum_{y \in Y_m} f(y) \left(\int_{\{\rho_m(F_x)=y\}} \alpha(x) d\mu(x) \right),$$

which means that $T_1(\alpha)$ is the complex-conjugate of the projection to $L_0^2(Y_m)$ of the function

$$y \mapsto \frac{1}{v_m(y)} \int_{\{\rho_m(F_x)=y\}} \alpha(x) d\mu(x)$$

on Y_m . For $m = \{\ell\}$, this projection is obtained by subtracting the contribution of the constant function, i.e., subtracting the average over y : it is

$$\begin{aligned} y \mapsto & \frac{1}{v(y)} \int_{\{\rho_\ell(F_x)=y\}} \alpha(x) d\mu(x) - \sum_y \int_{\{\rho_\ell(F_x)=y\}} \alpha(x) d\mu(x) \\ & = \frac{1}{v(y)} \int_{\{\rho_\ell(F_x)=y\}} \alpha(x) d\mu(x) - \int_X \alpha(x) d\mu(x). \end{aligned}$$

In the case of counting measure and a uniform density v , this becomes the quantity

$$\sum_{\rho_\ell(F_x)=y} \alpha(x) - \frac{1}{|Y_\ell|} \sum_x \alpha(x)$$

(after multiplying by $v(y)$), which is a typical ‘error term’ appearing in sieve axioms.

To prove Proposition 2.3, we follow the most commonly used approach in the classical case, which is due to Gallagher and differs from Montgomery’s original version.

We need two lemmas to start. For $m \in S(\Lambda)$, $y \in Y_m$, an element φ of the basis \mathcal{B}_m , and a square-integrable function $\alpha \in L^2(X, \mu)$, we write

$$S(m, y) = \int_{\{\rho_m(F_x)=y\}} \alpha(x) d\mu(x),$$

and

$$S(\varphi) = \int_X \alpha(x) \varphi(\rho_m(F_x)) d\mu(x), \quad (2.5)$$

where the integrals are defined because $\mu(X) < +\infty$ by assumption. The first lemma is the following:

Lemma 2.6 *We have for all $\ell \in \Lambda$ the relation*

$$\sum_{\varphi \in \mathcal{B}_\ell^*} |S(\varphi)|^2 = \sum_{y \in Y_\ell} \frac{|S(\ell, y)|^2}{v(y)} - \left| \int_X \alpha(x) d\mu(x) \right|^2.$$

Proof Expanding the square by Fubini's Theorem, the left-hand side is

$$\int_X \int_X \alpha(x) \overline{\alpha(y)} \sum_{\varphi \in \mathcal{B}_\ell^*} \varphi(\rho_\ell(F_x)) \overline{\varphi(\rho_\ell(F_y))} d\mu(x) d\mu(y).$$

Since $(\varphi)_{\varphi \in \mathcal{B}_\ell}$ is an orthonormal basis of the space of functions on Y_ℓ , we can expand the delta function $z \mapsto \delta(y, z)$ in this basis for fixed $y \in Y_\ell$. Since

$$\langle \delta(y, \cdot), \varphi \rangle = v(y) \varphi(y),$$

this expansion is equivalent with the identity

$$\sum_{\varphi \in \mathcal{B}_\ell} \varphi(y) \overline{\varphi(z)} = \frac{1}{v(y)} \delta(y, z). \quad (2.6)$$

Taking on the right-hand side the contribution of the constant function 1, we get in particular

$$\sum_{\varphi \in \mathcal{B}_\ell^*} \varphi(\rho_\ell(F_x)) \overline{\varphi(\rho_\ell(F_y))} = \frac{1}{v(\rho_\ell(F_x))} \delta(\rho_\ell(F_x), \rho_\ell(F_y)) - 1.$$

Inserting this in the first relation, we obtain

$$\begin{aligned}
\sum_{\varphi \in \mathcal{B}_\ell^*} |S(\varphi)|^2 &= \int \int_{\{\rho_\ell(F_x) = \rho_\ell(F_y)\}} \frac{\alpha(x) \overline{\alpha(y)}}{v(\rho_\ell(F_x))} d\mu(x) d\mu(y) \\
&\quad - \int_X \int_X \alpha(x) \overline{\alpha(y)} d\mu(x) d\mu(y) \\
&= \sum_{z \in Y_\ell} \frac{1}{v(z)} \int \int_{\{\rho_\ell(F_x) = z = \rho_\ell(F_y)\}} \alpha(x) \overline{\alpha(y)} d\mu(x) d\mu(y) \\
&\quad - \left| \int_X \alpha(x) d\mu(x) \right|^2 \\
&= \sum_{z \in Y_\ell} \frac{|S(\ell, z)|^2}{v(z)} - \left| \int_X \alpha(x) d\mu(x) \right|^2,
\end{aligned}$$

as desired. □

Here is the next lemma.

Lemma 2.7 *Let $(\Psi, \Upsilon, \Omega, \mathcal{L}^*)$ be as above. For any square-integrable function $x \mapsto \alpha(x)$ on X supported on the sifted set $S(X, \Omega; \mathcal{L}^*) \subset X$, and for any $m \subset \mathcal{L}^*$, we have*

$$\sum_{\varphi \in \mathcal{B}_m^*} |S(\varphi)|^2 \geq \left| \int_X \alpha(x) d\mu(x) \right|^2 \prod_{\ell \in m} \frac{v(\Omega_\ell)}{v(Y_\ell - \Omega_\ell)},$$

where $S(\varphi)$ is given by (2.5).

Proof As in the classical case (see, e.g., [67, Lemma 7.15]), the proof proceeds by induction on the number of elements in m . If $m = \emptyset$, the inequality is trivial (there is equality, in fact). If $m = \{\ell\}$ with $\ell \in \Lambda$ (in the arithmetic case, m is a prime), then $\ell \in \mathcal{L}^*$. Using Cauchy's inequality and the definition of the sifted set with the assumption on $\alpha(x)$ to restrict the support of integration to elements where $\rho_\ell(F_x) \notin \Omega_\ell$, we obtain:

$$\begin{aligned}
\left| \int_X \alpha(x) d\mu(x) \right|^2 &= \left| \sum_{\substack{y \in Y_\ell \\ y \notin \Omega_\ell}} S(\ell, y) \right|^2 \leq \left(\sum_{y \notin \Omega_\ell} v(y) \right) \left(\sum_{y \in Y_\ell} \frac{|S(\ell, y)|^2}{v(y)} \right) \\
&= v(Y_\ell - \Omega_\ell) \sum_{y \in Y_\ell} \frac{|S(\ell, y)|^2}{v(y)} \\
&= v(Y_\ell - \Omega_\ell) \left\{ \sum_{\varphi \in \mathcal{B}_\ell^*} |S(\varphi)|^2 + \left| \int_X \alpha(x) d\mu(x) \right|^2 \right\}
\end{aligned}$$

(by Lemma 2.6), hence the result by moving $\left| \int \alpha(x) d\mu \right|^2$ on the left-hand side, since $v(Y_\ell) = 1$.

The induction step is now immediate, relying on the fact that the function α is arbitrary and the sets \mathcal{B}_m^* are ‘multiplicative’: for $m \subset \mathcal{L}^*$, not a singleton, write $m = m_1 m_2 = m_1 \cup m_2$ with m_1 and m_2 non-empty (and still subsets of \mathcal{L}^*). Then we have

$$\sum_{\varphi \in \mathcal{B}_{m_1 m_2}^*} |S(\varphi)|^2 = \sum_{\varphi_1 \in \mathcal{B}_{m_1}^*} \sum_{\varphi_2 \in \mathcal{B}_{m_2}^*} |S(\varphi_1 \otimes \varphi_2)|^2$$

where $\varphi_1 \otimes \varphi_2$ is the function $(y, z) \mapsto \varphi_1(y)\varphi_2(z)$. For fixed φ_1 , we can express the inner sum as

$$S(\varphi_1 \otimes \varphi_2) = \int_X \beta(x) \varphi_2(\rho_{m_2}(F_x)) d\mu(x)$$

with $\beta(x) = \alpha(x)\varphi_1(\rho_{m_1}(F_x))$, which is also supported on $S(X, \Omega; \mathcal{L}^*)$. By the induction hypothesis applied first to m_2 , then to m_1 , we obtain

$$\begin{aligned} \sum_{\varphi \in \mathcal{B}_{m_1 m_2}^*} |S(\varphi)|^2 &\geq \prod_{\ell|m_2} \frac{v(\Omega_\ell)}{v(Y_\ell - \Omega_\ell)} \sum_{\varphi_1 \in \mathcal{B}_{m_1}^*} \left| \int_X \beta(x) d\mu(x) \right|^2 \\ &= \prod_{\ell|m_2} \frac{v(\Omega_\ell)}{v(Y_\ell - \Omega_\ell)} \sum_{\varphi_1 \in \mathcal{B}_{m_1}^*} |S(\varphi_1)|^2 \\ &\geq \prod_{\ell|m_1 m_2} \frac{v(\Omega_\ell)}{v(Y_\ell - \Omega_\ell)} \left| \int_X \alpha(x) d\mu(x) \right|^2. \end{aligned}$$

□

Now the proof of Proposition 2.3 is easy.

Proof of Proposition 2.3 Take $\alpha(x)$ to be the characteristic function of $S(X, \Omega; \mathcal{L}^*)$ and sum over $m \in \mathcal{L}$ the inequality of Lemma 2.7; since

$$\int_X \alpha(x) d\mu(x) = \int_X \alpha(x)^2 d\mu(x) = |S(X, \Omega; \mathcal{L}^*)|,$$

it follows that

$$|S(X, \Omega; \mathcal{L}^*)|^2 \sum_{m \in \mathcal{L}} \prod_{\ell|m} \frac{v(\Omega_\ell)}{v(Y_\ell - \Omega_\ell)} \leq \sum_{m \in \mathcal{L}} \sum_{\varphi \in \mathcal{B}_m^*} |S(\varphi)|^2 \leq \Delta |S(X, \Omega; \mathcal{L}^*)|,$$

hence the result. □

2.3 Duality and ‘exponential sums’

At this point a ‘large sieve inequality’ will be an estimate for the quantity Δ . There are various techniques available for this purpose, and we refer to [67, Chapter VII] for a survey of some of them. The simplest is the familiar duality principle for bilinear forms or linear operators. Since Δ is the square of the norm of a linear operator, it is the square of the norm of its adjoint. Hence we deduce:

Lemma 2.8 *Let $\Psi = (Y, \Lambda, (\rho_\ell))$ be a sieve setting, (X, μ, F) a siftable set, \mathcal{L} a sieve support associated to \mathcal{L}^* . Fix orthonormal bases \mathcal{B}_ℓ and define \mathcal{B}_m as above. Then the large sieve constant $\Delta(X, \mathcal{L})$ is the smallest number Δ such that*

$$\int_X \left| \sum_{m \in \mathcal{L}} \sum_{\varphi \in \mathcal{B}_m^*} \beta(m, \varphi) \varphi(\rho_m(F_x)) \right|^2 d\mu(x) \leq \Delta \sum_m \sum_{\varphi} |\beta(m, \varphi)|^2 \quad (2.7)$$

for all vectors of complex numbers $(\beta(m, \varphi))$.

The usefulness of this is that it leads to a bound for Δ in terms of estimates for the ‘dual’ sums $W(\varphi, \varphi')$ obtained by expanding the square in this inequality, i.e.,

$$W(\varphi, \varphi') = \int_X \varphi(\rho_m(F_x)) \overline{\varphi'(\rho_n(F_x))} d\mu(x),$$

where $\varphi \in \mathcal{B}_m$ and $\varphi' \in \mathcal{B}_n$ for some m and n in $S(\Lambda)$. Precisely, we have:

Proposition 2.9 *Let $\Psi = (Y, \Lambda, (\rho_\ell))$ be a sieve setting, $\Upsilon = (X, \mu, F)$ a siftable set, \mathcal{L}^* a prime sieve support and \mathcal{L} an associated sieve support. Then the large sieve constant satisfies*

$$\Delta \leq \max_{m \in \mathcal{L}} \max_{\varphi \in \mathcal{B}_m^*} \sum_{n \in \mathcal{L}} \sum_{\varphi' \in \mathcal{B}_n^*} |W(\varphi, \varphi')|. \quad (2.8)$$

Proof Expanding the left-hand side of (2.7), we have

$$\begin{aligned} & \int_X \left| \sum_{m \in \mathcal{L}} \sum_{\varphi \in \mathcal{B}_m^*} \beta(m, \varphi) \varphi(\rho_m(F_x)) \right|^2 d\mu(x) \\ &= \sum_{m, n} \sum_{\varphi, \varphi'} \beta(m, \varphi) \overline{\beta(n, \varphi')} W(\varphi, \varphi') \end{aligned}$$

and applying $|uv| \leq \frac{1}{2}(|u|^2 + |v|^2)$, the result follows directly. \square

The point is that sieve results are now reduced to individual *uniform* estimates for the ‘sums’ $W(\varphi, \varphi')$. Note that, here, the choice of the orthonormal basis may well be very important in estimating $W(\varphi, \varphi')$ and therefore Δ .

Remark 2.10 For some applications, it is useful to gain some analytic flexibility by introducing a smoothing factor. Abstractly, this would usually mean that $X \subset X'$ for some other set X' , μ is the restriction of a measure (still denoted μ) on X' and $x \mapsto F_x$ extends in some way to $X' \rightarrow Y$. Then for an arbitrary (integrable) function $\Phi : X' \rightarrow [0, 1]$ such that $\Phi(x) \geq 1$ for $x \in X$, and $\pi \in \Pi_m$, $\tau \in \Pi_n$, we consider the ‘smoothed’ sum

$$W_\Phi(\varphi, \varphi') = \int_{X'} \varphi(\rho_m(F_x)) \overline{\varphi'(\rho_n(F_x))} \Phi(x) d\mu(x);$$

then the large sieve constant $\Delta(X, \mathcal{L})$ satisfies

$$\Delta \leq \max_{m \in \mathcal{L}} \max_{\varphi \in \mathcal{B}_m^*} \sum_{n \in \mathcal{L}} \sum_{\varphi' \in \mathcal{B}_n^*} |W_\Phi(\varphi, \varphi')|.$$

Typically, take the case where $X = \{1, \dots, N\}$ for some N , let $X' = \{n \geq 1\}$ and let Φ be a smooth compactly supported function on $[0, +\infty[$ such that $0 \leq \Phi(x) \leq 1$, $\Phi(x) = 1$ for $0 \leq x \leq 1$ and $\Phi(x) = 0$ for $x \geq 2$. Then W_Φ is a typical ‘smoothed’ sum. These are useful (for instance) for the purposes of Mellin inversion or Fourier analysis, because the Mellin transform of Φ is holomorphic for $\text{Re}(s) > 0$ with fast decay in vertical strips (see the proof of Proposition G.3 for an example of use of this technique; the smoothness of Φ is what translates to fast decay, whereas a discontinuous function such as the characteristic function of an interval has much worse behaviour).

We do not need to introduce and keep track of this generalization however, since we can simply obtain the desired result by using the siftable set $(X', \Phi\mu, F)$ instead of (X, μ, F) and the inequality

$$|S(X, \Omega; \mathcal{L}^*)| \leq |S(X', \Omega; \mathcal{L}^*)|.$$

Note that the above ‘smoothed’ bound for Δ (which is indeed correct) is the same as (2.8) for the new siftable set. All this is of course an indication that the generality we work in is indeed useful.

So we come back to the study of the general sums $W(\varphi, \varphi')$. At least formally, we can proceed in full generality as follows, where the idea is that in applications $\rho_m(F_x)$ should range fairly equitably (with respect to the density v_m) over the elements of Y_m , so the sum $W(\varphi, \varphi')$ should be estimated by exploiting the fact that the values $\varphi(\rho_m(F_x))\overline{\varphi'(\rho_n(F_x))}$ depend only on $\rho_m(F_x)$ and $\rho_n(F_x)$. To do this, we introduce further notation.

Let m, n be two elements of $S(\Lambda)$, $\varphi \in \mathcal{B}_m$, $\varphi' \in \mathcal{B}_n$. Let $d = m \cap n$ be the intersection (g.c.d. in the case of integers) of m and n , and write $m = m'd = m' \cup d$, $n = n'd = n' \cup d$ (disjoint unions). According to the multiplicative definition of \mathcal{B}_m and \mathcal{B}_n , we can write

$$\varphi = \varphi_{m'} \otimes \varphi_d, \quad \varphi' = \varphi'_{n'} \otimes \varphi'_d$$

for some unique basis elements $\varphi_{m'} \in \mathcal{B}_{m'}$, $\varphi_d, \varphi'_d \in \mathcal{B}_d$ and $\varphi'_{n'} \in \mathcal{B}_{n'}$.

Let $[m, n] = m \cup n$ be the ‘l.c.m.’ of m and n . We have the decomposition

$$Y_{[m,n]} = Y_{m'} \times Y_d \times Y_{n'},$$

the (not necessarily surjective) map $\rho_{[m,n]} : Y \rightarrow Y_{[m,n]}$ and the function

$$[\varphi, \overline{\varphi'}] = \varphi_{m'} \otimes (\varphi_d \overline{\varphi'_d}) \otimes \overline{\varphi'_{n'}} : (y_1, y_d, y_2) \mapsto \varphi_{m'}(y_1) \varphi_d(y_d) \overline{\varphi'_d(y_d) \varphi'_{n'}(y_2)} \quad (2.9)$$

(which, usually, is not a basis element in $\mathcal{B}_{[m,n]}$).

The motivation for all this is the following tautology:

Lemma 2.11 *Let m, n, φ and φ' be as before. We have*

$$[\varphi, \overline{\varphi'}](\rho_{[m,n]}(y)) = \varphi(\rho_m(y)) \overline{\varphi'(\rho_n(y))}$$

for all $y \in Y$, hence

$$W(\varphi, \varphi') = \int_X [\varphi, \overline{\varphi'}](\rho_{[m,n]}(F_x)) d\mu(x).$$

Now we can hope to split the integral according to the value of $y = \rho_{[m,n]}(F_x)$ in $Y_{[m,n]}$, and evaluate it by first summing the main term in an equidistribution statement. More precisely, for $d \in S(\Lambda)$ and $y \in Y_d$, we define $r_d(X; y)$ as the ‘error term’ in the expected equidistribution statement:

$$|\{\rho_d(F_x) = y\}| = \int_{\{\rho_d(F_x)=y\}} d\mu(x) = v_d(y)|X| + r_d(X; y). \quad (2.10)$$

This, and what follows, only makes sense if $r_d(X; y)$ is of smaller order of magnitude than the main term. For fixed d , such is the case precisely if we have a sequence of siftable sets (X_n, μ_n, F_n) such that the image measures $(\rho_d \circ F_n)_*(\mu_n)$ on Y_d converge weakly to the measure $|X|v_d$, or in other words, if $\rho_d(F_{n,x})$ is equidistributed with respect to this measure. It is in this sense that there is no real choice of v_d : in order for the large sieve principle to apply efficiently, this type of individual (for one d) equidistribution is necessary and fixes the density v_d .

Having defined $r_d(X; y)$, we can compute $W(\varphi, \varphi')$ as sketched:

$$\begin{aligned}
 W(\varphi, \varphi') &= \int_X [\varphi, \overline{\varphi'}](\rho_{[m,n]}(F_x)) d\mu(x) \\
 &= \sum_{y \in Y_{[m,n]}} [\varphi, \overline{\varphi'}](y) \int_{\{\rho_{[m,n]}(F_x) = y\}} d\mu(x) \\
 &= m([\varphi, \overline{\varphi'}]|X| + O\left(\sum_{y \in Y_{[m,n]}} \|\varphi, \overline{\varphi'}\|_\infty |r_{[m,n]}(X; y)|\right) \quad (2.11)
 \end{aligned}$$

after inserting (2.10), where the implied constant is of modulus ≤ 1 and

$$m([\varphi, \overline{\varphi'}]) = \sum_{y \in Y_{[m,n]}} v_{[m,n]}(y) [\varphi, \overline{\varphi'}](y) = \langle [\varphi, \overline{\varphi'}], 1 \rangle,$$

the inner product in $L^2(Y_{[m,n]})$. This is easy to evaluate:

Lemma 2.12 *With notation as before, we have*

$$m([\varphi, \overline{\varphi'}]) = \delta((m, \varphi), (n, \varphi')).$$

Proof From the definition of the inner products on Y_d , $d \in S(\Lambda)$, and (2.9), we have

$$m([\varphi, \overline{\varphi'}]) = \langle \varphi_{m'}, 1 \rangle_{Y_{m'}} \langle \varphi_d, \varphi'_d \rangle_{Y_d} \langle 1, \varphi'_{n'} \rangle_{Y_{n'}}.$$

The two extreme terms are zero, unless $m' = n' = \emptyset$ (i.e., $m' = n' = 1$ in the case of squarefree integers). In this case, we have $m = n = d$, $\varphi = \varphi_d$, $\varphi' = \varphi'_d$, and then the middle term is $\delta(\varphi, \varphi')$ by orthonormality. \square

Hence we deduce the reduction of the large sieve to equidistribution:

Corollary 2.13 *Let $(Y, \Lambda, (\rho_\ell))$ be a sieve setting, (X, μ, F) a siftable set. Define the equidistribution remainder terms $r_d(X; y)$ for $d \in S(\Lambda)$ and $y \in Y_d$ by (2.10). Then for any sieve support \mathcal{L} , the large sieve constant $\Delta(X, \mathcal{L})$ is bounded by*

$$\Delta \leq |X| + \max_{m, \varphi} \sum_n \left(\sum_{y \in Y_{[m,n]}} |r_{[m,n]}(X; y)| \right) \left(\sum_{\varphi'} \|\varphi, \overline{\varphi'}\|_\infty \right).$$

In general, the bound arising from this corollary is quite a bit weaker than the more natural one arising from a direct study of the sums $W(\varphi, \varphi')$. However, it at least provides a measure of sieve whenever a quantitative equidistribution statement is known, and qualitatively, it may be comparable.

Remark 2.14 The ‘equidistribution’ approach and the exponential sums technique are very closely related. Indeed, in the opposite direction (from $W(\varphi, \varphi')$ to equidistribution), given a subset $A \subset Y_d$, we can expand the characteristic function χ_A of A in terms of the basis \mathcal{B}_d , and write

$$\chi_A(\rho_d(F_x)) = \sum_{\varphi \in \mathcal{B}_d} \langle \chi_A, \varphi \rangle \varphi(\rho_d(F_x)),$$

from which one gets

$$|\{\rho_d(F_x) \in A\}| = v_d(A)|X| + \sum_{\varphi \in \mathcal{B}_d - \{1\}} \langle \chi_A, \varphi \rangle W(1, \varphi)$$

(so the remainder term in (2.10) is itself a combination of exponential sums). In particular, applying now Cauchy’s inequality, the fact that \mathcal{B}_d is an orthonormal basis of $L^2(Y_d)$, and that $\|\chi_A\| = \sqrt{v_d(\Omega)}$, we derive

$$|\{\rho_d(F_x) \in A\}| = v_d(A)|X| + O\left(\sqrt{v_d(A)} \left(\sum_{\varphi \in \mathcal{B}_d - \{1\}} W(1, \varphi)^2\right)^{1/2}\right) \quad (2.12)$$

where the implied constant is ≤ 1 . We will use this later on in some cases, but note that it may also be seen as another instance of the large sieve principle, though fairly degenerate: it amounts to taking the sieve setting $(Y, \{d\}, \rho_d)$ while keeping the original siftable set, and choosing Ω_d to be the complement of A (since we looked at $x \in X$ with $\rho_d(F_x)$ in A).

In all cases considered in this book (except the simplest), equidistribution is proved in this manner, and then one might as well deduce the large sieve constant from the bounds for general sums $W(\varphi, \varphi')$ – not doing so means, essentially, performing the same operation forward and backward and weakening the estimates in both directions . . .

2.4 The dual sieve

The equivalent definition of the large sieve constant by means of the duality principle (i.e., Lemma 2.8) is quite useful in itself. For instance, it yields the following type of sieve inequality, which in the classical case goes back to Rényi.

Proposition 2.15 *Let $(Y, \Lambda, (\rho_\ell))$ be a sieve setting, (X, μ, F) a siftable set and \mathcal{L}^* a prime sieve support. Let Δ be the large sieve constant for $\mathcal{L} = \mathcal{L}^{*,2}$*

² Precisely, \mathcal{L} is the set of singletons $\{\ell\}$ for $\ell \in \mathcal{L}^*$.

Then for any sieving sets (Ω_ℓ) , we have

$$\int_X (P(x, \mathcal{L}) - P(\mathcal{L}))^2 d\mu(x) \leq \Delta Q(\mathcal{L}) \quad (2.13)$$

where

$$P(x, \mathcal{L}) = \sum_{\substack{\ell \in \mathcal{L} \\ \rho_\ell(F_x) \in \Omega_\ell}} 1, \quad P(\mathcal{L}) = \sum_{\ell \in \mathcal{L}} v(\Omega_\ell), \quad (2.14)$$

$$Q(\mathcal{L}) = \sum_{\ell \in \mathcal{L}} v(\Omega_\ell)(1 - v(\Omega_\ell)). \quad (2.15)$$

Proof By expanding the characteristic function χ_ℓ of $\Omega_\ell \subset Y_\ell$ in the orthonormal basis \mathcal{B}_ℓ , we obtain

$$P(x, \mathcal{L}) = P(\mathcal{L}) + \sum_{\ell \in \mathcal{L}} \sum_{\varphi \in \mathcal{B}_\ell^*} \beta(\ell, \varphi) \varphi(\rho_\ell(F_x)),$$

where

$$\beta(\ell, \varphi) = \sum_{y \in \Omega_\ell} v_\ell(y) \overline{\varphi}(y),$$

and we used the fact that $\mathcal{B}_\ell^* = \mathcal{B}_\ell - \{1\}$ for $\ell \in \Lambda$. Thus we get

$$\begin{aligned} \int_X (P(x, \mathcal{L}) - P(\mathcal{L}))^2 d\mu(x) &= \int_X \left| \sum_{\ell \in \mathcal{L}} \sum_{\varphi \in \mathcal{B}_\ell^*} \beta(\ell, \varphi) \varphi(\rho_\ell(F_x)) \right|^2 d\mu(x) \\ &\leq \Delta \sum_{\ell \in \mathcal{L}} \sum_{\varphi \in \mathcal{B}_\ell^*} |\beta(\ell, \varphi)|^2 \end{aligned}$$

by applying (2.7). Since we have

$$\begin{aligned} \sum_{\varphi \in \mathcal{B}_\ell^*} |\beta(\ell, \varphi)|^2 &= \sum_{\varphi \in \mathcal{B}_\ell} |\beta(\ell, \varphi)|^2 - |\beta(\ell, 1)|^2 \\ &= \|\chi_\ell\|^2 - v(\Omega_\ell)^2 = v(\Omega_\ell)(1 - v(\Omega_\ell)), \end{aligned}$$

this implies the result. \square

In particular, since $P(x, \mathcal{L}) = 0$ for $x \in S(X, \Omega; \mathcal{L}^*)$ and $Q(\mathcal{L}) \leq P(\mathcal{L})$, we get (by positivity) the estimate

$$|S(X, \Omega; \mathcal{L}^*)| \leq \Delta P(\mathcal{L})^{-1},$$

which is the analogue of the inequalities used, e.g., by Gallagher in [46, Theorem A], and by the author in [80]. This inequality also follows from Proposition 2.3 if we take \mathcal{L} containing only singletons (in the arithmetic case,

this means using only the primes), since we get the estimate

$$|S(X, \Omega; \mathcal{L}^*)| \leq \Delta H^{-1}$$

with

$$H = \sum_{\ell \in \mathcal{L}} \frac{v(\Omega_\ell)}{v(Y_\ell - \Omega_\ell)} \geq \sum_{\ell \in \mathcal{L}} v(\Omega_\ell) = P(\mathcal{L})$$

(in fact, by Cauchy's inequality, we have $P(\mathcal{L})^2 \leq H Q(\mathcal{L})$).

This type of result is also related to Turán's method in probabilistic number theory. If we try to use it to count primes (or twin primes, for instance, or more generally if we look at small sieve situations), it seems quite weak.

Indeed, taking X to be the set of positive integers $\leq N$, \mathcal{L}^* the set of primes $\leq L$ (with the large sieve constant $\Delta = N - 1 + L^2$ which comes from the classical large sieve inequality, see (4.1)), we get

$$\sum_{n \leq N} \left(\omega_L(n) - \sum_{\ell \leq L} \frac{1}{\ell} \right)^2 \leq (N - 1 + L^2) \sum_{\ell \leq L} \frac{1}{\ell}, \quad (2.16)$$

when $\Omega_\ell = \{0\}$, where $\omega_L(n)$ is the number of prime divisors of n which are $\leq L$. In turn, since

$$\sum_{\ell \leq L} \frac{1}{\ell} = \log \log L + O(1)$$

for $L \geq 2$, this only implies that

$$\pi(N) \ll \frac{N}{\log \log N},$$

if we want to estimate the number of primes $\leq N$.

However, the dual sieve inequality is really a different type of statement, and it carries some useful additional flexibility: for instance, it still implies that for $n \geq 3$ we have

$$|\{n \leq n \mid \omega(n) < \kappa \log \log N\}| \ll \frac{N}{\log \log N}$$

for any $\kappa \in]0, 1[$, the implied constant depending only on κ .

Moreover, the estimate (2.16) is *of the right order of magnitude* for $L = N^{1/2}$. Indeed, Turán proved that

$$\sum_{n \leq N} (\omega(n) - \log \log N)^2 = N \log \log N + O(1) \quad (2.17)$$

for $N \geq 2$ (see, e.g., [99, Exercise 2.3.8]).

In addition, according to the Erdős–Kac theorem, we have

$$\frac{1}{N} \left| \left\{ n \leq N \mid \alpha \leq \frac{\omega(n) - \log \log N}{\sqrt{\log \log N}} \leq \beta \right\} \right| \rightarrow \frac{1}{\sqrt{2\pi}} \int_{\alpha}^{\beta} e^{-t^2/2} dt$$

as $N \rightarrow +\infty$, for any fixed $\alpha, \beta \in \mathbf{R}$. In other words, for large N , $\omega(n)$ behaves on average over $n \leq N$ like a normal random variable with mean $\log \log N$ and variance $\sqrt{\log \log N}$; this is related to the Central Limit Theorem, see Appendix F, and [10] for a probabilistic explanation.

All this shows in particular that one can not hope to improve (2.13) in general by using information related to all ‘squarefree’ numbers.

These remarks indicate clearly that Proposition 2.15 is of independent interest in cases where the ‘stronger’ form of the large sieve is in fact not adapted to the type of situation considered. In Section 4.4, we will describe an amusing use of the inequality (2.13), where the ‘pure sieve’ bound would indeed be essentially trivial, and in Chapters 6 and 7, we will use it to get some results on the number of prime divisors of integers constructed in rather exotic ways . . .

2.5 General comments on the large sieve inequality

This rather philosophical section will attempt to explain the meaning of the large sieve principle, and in particular what can be expected from it. Readers already familiar with sieve methods can probably go to the next chapter.

We assume that the sieve setting $(Y, \Lambda, (\rho_{\ell}))$ is fixed. Two quantities arise in the sieve bound, and must be dealt with before it may be successfully applied: the large sieve constant Δ , which depends on X and \mathcal{L} , but not on the sieving sets, and the saving factor H , which depends on \mathcal{L} and on the sieving sets, but not on X . This ‘separation of variables’ is one of the keys to the usefulness of the sieve. Knowing a bound for Δ , many sieves are reduced to evaluations of H , and similarly, if we know how to evaluate H for certain types of sieving sets, we can attack the study of Δ knowing that many applications will arise. From this, the vaunted *uniformity of sieve methods*, which is one of the main explanations for their power in number-theoretic applications, arises. Indeed, other types of asymptotic counting methods (e.g., generating series and tauberian arguments of one form or another) often have very poor uniformity. The best example of this situation is the distribution of primes in arithmetic progressions, where sieve methods quickly yield the Brun–Titchmarsh inequality (see (6.3)), which goes well beyond even the reach of the Generalized Riemann Hypothesis, and in comparison to which the best unconditional result, the Siegel–Walfisz Theorem,

is pitifully weak in its uniformity. On the other hand, of course, sieve is usually constrained to yield inequalities only, whereas other methods can provide asymptotic formulas, often with strong error terms.

We start by discussing the saving factor H , which is not where we put the main emphasis in this book. Notice that the dependency of H on the sieving sets Ω_ℓ involves only the density $\nu_\ell(\Omega_\ell)$, not the specific structure of Ω_ℓ . This is one of the sources of the uniformity of sieves, because it means that very different-looking problems can be reduced to the same computation, and indeed to computations which have already been done. On the other hand, this clearly limits how far the sieve can go, since one may expect that the true value of the size of the sifted set depends on more than the ‘local’ densities. This seems especially true in situations which are genuinely of large sieve type (which we take to mean that $\nu(\Omega_\ell)$ is bounded from below), and may well be the reason why the best bounds for the number of integral monic polynomials with ‘small’ splitting fields (see [46], and Chapter 4) remain rather far from the expected truth.³ There is almost nothing known about this issue. However, very recently, H. Helfgott and A. Venkatesh [60] have proved a remarkable result that gives the first insight about the phenomena that may occur. Roughly speaking, they show that when considering the two-dimensional sieve setting

$$(\mathbf{Z}^2, \{\text{primes}\}, \mathbf{Z}^2 \rightarrow (\mathbf{Z}/\ell\mathbf{Z})^2)$$

with $X = \{(n, m) \mid 0 \leq n, m \leq N\}$ for some $N \geq 1$ (with counting measure and $F_x = x$), if one has a *sifted* set $S \subset X$, such that the subsets of permitted residues classes (those which intersect S modulo ℓ , roughly the complement of Ω_ℓ) are of size $\leq \kappa \ell$ for some $\kappa > 0$ and all primes ℓ , then either the set S is extremely small ($|S| \ll N^\varepsilon$ for any $\varepsilon > 0$), or there exists a plane algebraic curve of bounded degree (in terms of κ and ε) which contains at least $(1 - \varepsilon)|S|$ elements of S . This can be thought of as an additive combinatorics dichotomy: a set is either very random (here, very small), or has structure (here, is ‘almost’ algebraic). This result is likely to have a strong influence on the finer (finest!) development of the large sieve in the future, but we will not go into such directions at all in this book. (Note that the methods of Helfgott and Venkatesh are elementary and related to the so-called ‘larger sieve’ of Gallagher.)

The form of H may seem surprising at first. However, its nature becomes clearer if one takes for sieve support \mathcal{L} , the full power-set of \mathcal{L}^* . In fact, we then find (reverting to inclusion notation for clarity) that

³ This is relative; in other situations involving irreducible polynomials, we would be very happy to obtain such a good bound as Gallagher’s!

$$H = \sum_{m \subset \mathcal{L}^*} \prod_{\ell \in m} \frac{v(\Omega_\ell)}{1 - v(\Omega_\ell)} = \prod_{\ell \in m} \left(1 + \frac{v(\Omega_\ell)}{1 - v(\Omega_\ell)} \right) = \prod_{\ell \in m} \frac{1}{1 - v(\Omega_\ell)}$$

and hence H^{-1} is the product

$$\prod_{\ell \in \mathcal{L}^*} v_\ell(Y_\ell - \Omega_\ell),$$

which is equal to the probability (for the product measure on $Y_{\mathcal{L}^*}$) that a random element $(y_\ell)_{\ell \in \mathcal{L}^*}$ has components outside of Ω_ℓ for all $\ell \in \mathcal{L}^*$. Under an assumption of equidistribution of the images F_x of $x \in X$ under the maps ρ_ℓ , and of independence of the various ℓ 's, this is the expected density of the sifted set. Hence, one should see H , in the general case when \mathcal{L} is *not* the whole set of subsets of \mathcal{L}^* , as an approximation to this ideal density. The point is that we can not hope to control the large sieve constant Δ for such a large sieve support (which has exponentially many elements compared to \mathcal{L}^*), and hence in practice we reduce (drastically) the size of \mathcal{L} in order to make Δ manageable, while it is possible to show that the size of H does not decrease too much. In fact, some of the very first examples of sieve show that this type of trade-off is necessary (indeed, the number of primes $< x$ is $\sim x/\log x$ as $x \rightarrow +\infty$ by the Prime Number Theorem, and the expected density for integers $n < x$ not divisible by primes $< y$ is⁴

$$\prod_{p < y} (1 - p^{-1}) \sim e^{-\gamma} \frac{1}{\log y}$$

as $y \rightarrow +\infty$, so that the order of magnitude of the expected density is correct when $y = \sqrt{x}$, detecting primes, but not the leading term). So the bargain is a very good one.

When the index set Λ is a subset of the set of prime numbers, which is the case in almost all applications we know, evaluating H typically boils down to the well-understood general problem of finding lower bounds for sums

$$\sum_{m \leq L}^b f(m) = \sum_{m \leq L} \mu(m)^2 f(m)$$

where $f(m)$ is a multiplicative function of m , i.e., one such that $f(1) = 1$ and $f(ab) = f(a)f(b)$ when a and b are coprime (note that $\mu(m)^2 f(m)$ is also multiplicative). This problem is not a trivial one, of course (as anyone not

⁴ This is the Mertens formula, where γ is Euler's constant, see, e.g. [99, Theorem 2.7].

already acquainted with the results and techniques should convince themselves by trying to find the asymptotic behaviour of, say,

$$\sum_{m \leq L}^b \frac{3^{\omega(m)} \varphi(m)}{\psi(m)} \prod_{\substack{\ell | m \\ \ell \equiv 1 \pmod{4}}} \left(1 + \frac{3}{\ell}\right); \quad (2.18)$$

see Exercise G.1 for the solution). Yet, in the cases which naturally occur in sieve theory (classical and otherwise), there is an extensive literature available,⁵ and we can select for our applications very strong results of various types. Even if we select a sieve support \mathcal{L} other than the traditional one of squarefree integers $\leq L$, as we will do at some point (and as Zywinia also did) with

$$\mathcal{L} = \{m \mid m \text{ is squarefree and } g(m) \leq L\}$$

where g is some other multiplicative function ‘close to m ’ on average, bounds for

$$\sum_{g(m) \leq M}^b f(m)$$

are also known (we will use a very general result of Lau and Wu [88]). We give a short survey of some of those estimates in Appendix G.

In fact, from the point of view of the new applications considered here, it is the computation of the densities $\nu(\Omega_\ell)$ themselves which turns out to be sometimes quite deep and delicate. Indeed, when Ω_ℓ is a subset of a matrix group over a finite field, as will be the case in Chapters 7 and 8, we need to appeal to non-trivial structure results on such groups, due to A. Borel and N. Chavdarov. These will be summarized in Appendix B.

This being said, from now on we look at the large sieve constant and at the inequality

$$\sum_{m \in \mathcal{L}} \sum_{\varphi \in \mathcal{B}_m^*} \left| \int_X \alpha(x) \varphi(\rho_m(F_x)) d\mu(x) \right|^2 \leq \Delta \|\alpha\|^2 \quad (2.19)$$

that defines it, to indicate its meaning, partly from a more analytic point of view. In particular, we discuss what should be expected or hoped for, as to the value of Δ in a given sieve setting. Not only is this useful to understand the sieve itself, but it suggests the possibility of other applications arising directly from (2.19) (as turned out to be the case with the classical large sieve).

⁵ Among the many names that can be mentioned, we mention Wirsing, Selberg and Delange.

First of all, there is a trivial upper bound: applying the Cauchy–Schwarz inequality we have

$$\sum_{m \in \mathcal{L}} \sum_{\varphi \in \mathcal{B}_m^*} \left| \int_X \alpha(x) \varphi(\rho_m(F_x)) d\mu(x) \right|^2 \leq \left(\sum_{m \in \mathcal{L}} \sum_{\varphi \in \mathcal{B}_m^*} W(\varphi, \varphi) \right) \|\alpha\|^2.$$

In many cases the basis functions φ are bounded (on average over m at least), which suggests that $W(\varphi, \varphi)$ (the L^2 -norm of $x \mapsto \varphi(\rho_m(F_x))$) is of size comparable to $|X|$, leading to an upper bound of the form

$$\Delta \leq |X|N(\mathcal{L}),$$

where

$$N(\mathcal{L}) = \sum_{m \in \mathcal{L}} \sum_{\varphi \in \mathcal{B}_m^*} 1 = \sum_{m \in \mathcal{L}} \prod_{\ell | m} (|Y_\ell| - 1).$$

Not surprisingly, this trivial bound is useless for sieve purposes (or for any other application). Closer to the truth, at least in important cases, is the lower bound for Δ arising by choosing special functions α . Indeed, taking

$$\alpha(x) = \overline{\varphi(\rho_m(F_x))}$$

for some fixed m and $\varphi \in \mathcal{B}_m^*$, we find a lower bound that amounts to

$$\Delta \geq \max_{m, \varphi} W(\varphi, \varphi),$$

and again this should be of size roughly $|X|$.

Another way of seeing or phrasing this lower bound is obtained from the duality principle. First of all, although the quantity on the right-hand side of (2.8) certainly depends on the orthonormal basis, it is a fact that there always exists a basis for which the bound on the right is equal to Δ (at least if the bases are not constrained to be defined ‘multiplicatively’ from the prime sieve support \mathcal{L}^* ; one need only choose a basis which diagonalizes the self-adjoint operator T^*T , where T is defined in Remark 2.5).

In (2.8), for a given (n, φ) , the sum

$$\sum_{m \in \mathcal{L}} \sum_{\varphi' \in \mathcal{B}_m^*} |W(\varphi, \varphi')|$$

contains the ‘diagonal’ term $W(\varphi, \varphi)$, which is the lower bound we obtained previously, and is expected to be comparable with $|X|$. Thus this sum is unlikely to be smaller than $|X|$. The expectation is that if \mathcal{L}^* is not too big, the sum of the other terms is at most of the same order of magnitude. Typically, assume for definiteness that \mathcal{L}^* is the set of primes $\leq L$, \mathcal{L} the set of squarefree integers $\leq L$.

Then we may hope for an inequality of the form

$$\Delta \leq |X| + |X|^{1-\alpha} L^A$$

for some $\alpha > 0$, and some $A \geq 0$. If this is true, then the sieve inequality becomes

$$|S(X, \Omega; \mathcal{L}^*)| \leq \frac{|X| + |X|^{1-\alpha} L^A}{H},$$

to be compared with the trivial estimate $|S(X, \Omega; \mathcal{L}^*)| \leq |X|$. This means that as long as

$$L^A \leq |X|^{1-\alpha}, \quad H > 2,$$

the sieve bound is non-trivial. Since we often have a sequence of siftable sets where $|X| \rightarrow +\infty$, this a-priori bound allows L to grow also, and then usually $H \rightarrow +\infty$, giving a sizable gain on the trivial bound. An important point is that, qualitatively, the effect is mostly unchanged however small α is, and however big A is – at least for direct applications of the sieve.

If we look in turn to what is the best we can hope for, note that if we expand the square in (2.19) directly, we are led to sums

$$\tilde{W}(x, y) = \sum_m \sum_{\varphi \in \mathcal{B}_m^*} \varphi(\rho_m(F_x)) \overline{\varphi(\rho_m(F_y))}$$

for $x, y \in X$, dual to $W(\varphi, \varphi')$. We still have

$$\Delta \leq \max_x \int_X |\tilde{W}(x, y)| d\mu(y).$$

A good estimate for these sums is of the type

$$\tilde{W}(x, y) = \delta(x, y) \sum_m \sum_{\varphi \in \mathcal{B}_m^*} 1 + (\text{small remainder})$$

and this leads to the conclusion that Δ can not be expected to be smaller than $N(\mathcal{L})$. So the best possible outcome is that Δ be roughly of size $\max(|X|, N(\mathcal{L})) \asymp |X| + N(\mathcal{L})$. Note that the sums $W(\varphi, \varphi')$ are in fact usually simpler, often much simpler, to deal with; one reason being that, given $x \in X$, it may be hard to estimate directly the measure of the set of y where $F_y = F_x$, and $\tilde{W}(x, y) = \tilde{W}(x, x) > 0$ for any such y . (If F is injective, this is not an issue, but there may be other problems.)

The best possible bound is *indeed valid in the classical large sieve inequality*, see Theorem 4.1. This is now well known, but it may well be thought of as being surprising (there is no doubt it was considered an impressive and surprising result when first discovered). It seems rather unwise to hope for such a strong result in all situations. Indeed, suppose for simplicity that X is a finite set with counting measure. Another viewpoint coming from seeing $W(\varphi, \varphi')$ as an

exponential sum ('square-root cancellation') is that we can not expect better individual bounds than

$$W(\varphi, \varphi') \ll |X|^{1/2},$$

for $\varphi \in \mathcal{B}_m^*$, $\varphi' \in \mathcal{B}_n^*$ (disregarding the dependency on m, n , which is quite unrealistic), and even such optimistic assumptions only suggest the bound $|X| + |X|^{1/2}N(\mathcal{L})$.

Finally, let us come back to Corollary 2.13. We see that the bound for Δ is of the type we expect, provided the remainders $r_d(X; y)$ defined in (2.10) are rather smaller than $|X|$, at least on average over d and y in some range. We have already mentioned that this essentially fixes v_d . More importantly, since Y_d is the product of Y_ℓ for $\ell \mid d$, and v_d is the product measure of v_ℓ , the same must be true for the limiting distribution (when dealing with a sequence of siftable sets) of $\rho_d(F_x) = (\rho_\ell(F_x))_{\ell \mid d}$. This is an expression of asymptotic independence of the 'reductions' of the image of the map F . In particular, since $v_d(y) > 0$ by assumption, a necessary condition (beyond individual equidistribution for $\ell \in \Lambda$) is that the maps $\rho_d : Y \rightarrow Y_d$ all be surjective.

Definition 2.16 A sieve setting $\Psi = (Y, \Lambda, (\rho_\ell))$ is linearly disjoint if the map $\rho_m : Y \rightarrow Y_m$ is onto for all $m \in S(\Lambda)$.

In the simplest case where $Y = \mathbf{Z}$, Λ is the set of primes and ρ_ℓ is the reduction modulo ℓ , linear disjointness holds: this is simply the Chinese Remainder Theorem.

If a sieve setting is not linearly disjoint, this may well be because, in a sense, it has been badly chosen. Suppose, for instance, that there exists another set Z and maps $Y \xrightarrow{d} Z$ and $Y_\ell \xrightarrow{d_\ell} Z$ for which the triangles

$$\begin{array}{ccc} Y & \xrightarrow{\rho_\ell} & Y_\ell \\ & \searrow d & \swarrow d_\ell \\ & Z & \end{array}$$

commute. In this case, even though $Y \rightarrow Y_\ell$ may be surjective, we have

$$\rho_d(y) \in \{(x_\ell)_{\ell \mid d} \mid d_\ell(x_\ell) = d(y) \text{ for all } \ell \mid d\} \subset Y_d,$$

which implies that (except for trivial cases) ρ_d is not surjective if d has at least two elements. In a situation of this type, the natural step to take is to replace Y by the various sets $d^{-1}(z)$ for fixed $z \in Z$, Y_ℓ by $d_\ell^{-1}(z)$, which gives a 'second' chance of defining a linearly disjoint sieve setting . . .

3

Group and conjugacy sieves

We now come to the description of a more specific type of sieve setting, related to a group structure on Y . This exhausts most examples of applications we know at the moment.

3.1 Conjugacy sieves

A *group sieve* corresponds to a sieve setting $\Psi = (G, \Lambda, (\rho_\ell))$ where G is a group and the maps $\rho_\ell : G \rightarrow G_\ell$ are homomorphisms onto finite groups. A *conjugacy sieve*, similarly, is a sieve setting $\Psi = (G, \Lambda, (\rho_\ell))$ where $\rho_\ell : G \rightarrow G_\ell^\#$ is a surjective map from G to the finite set of conjugacy classes $G_\ell^\#$ of a finite group G_ℓ , that factors as

$$G \rightarrow G_\ell \rightarrow G_\ell^\#$$

where $G \rightarrow G_\ell$ is a group homomorphism (necessarily surjective since the image intersects every conjugacy class, so must be equal to G_ℓ ; this is a classical property of finite groups, which will play another role in one application of the sieve later on; see the proof of Theorem 4.2). Obviously, any group sieve induces a conjugacy sieve. Also, if G is abelian, group and conjugacy sieves are identical.

The group structure suggests a natural choice of orthonormal basis \mathcal{B}_ℓ for functions on G_ℓ or $G_\ell^\#$, as well as natural densities ν_ℓ . We start with the simpler conjugacy sieve.

From the classical representation theory of finite groups (see, e.g., [115], or Appendix C), we know that for any $\ell \in \Lambda$, the characters of G_ℓ , i.e., the functions

$$y \mapsto \text{Tr } \pi(y),$$

on G_ℓ , where π runs over the set Π_ℓ of (isomorphism classes of) irreducible linear representations $\pi : G_\ell \rightarrow GL(V_\pi)$, form an orthonormal basis of the space of functions on G_ℓ invariant under conjugation, with the inner product

$$\langle f, g \rangle = \frac{1}{|G_\ell|} \sum_{y \in G_\ell} f(y) \overline{g(y)}.$$

Translating this statement to functions on the set G_ℓ^\sharp of conjugacy classes, this means that the functions

$$\varphi(y^\sharp) = \text{Tr } \pi(y^\sharp)$$

on G_ℓ^\sharp form an orthonormal basis \mathcal{B}_ℓ of $L^2(G_\ell^\sharp)$ with the inner product

$$\langle f, g \rangle = \frac{1}{|G_\ell|} \sum_{y^\sharp \in G} |y^\sharp| f(y^\sharp) \overline{g(y^\sharp)}.$$

Moreover, the trivial representation 1 of G_ℓ has for character the constant function 1, so we can use the basis $\mathcal{B}_\ell = (\text{Tr } \pi(y^\sharp))_\pi$ for computing the large sieve constant if the density

$$v_\ell(y^\sharp) = \frac{|y^\sharp|}{|G_\ell|}$$

is used. Note that this is the image on G_ℓ^\sharp of the uniform density on G_ℓ .

Note also that in the abelian case, the representations are one-dimensional, and the basis thus described is the basis of characters of G_ℓ , with the uniform density, i.e., that of group homomorphisms $G_\ell \rightarrow \mathbf{C}^\times$ with

$$\langle f, g \rangle = \frac{1}{|G_\ell|} \sum_{y \in G_\ell} f(y) \overline{g(y)}.$$

Coming back to a general group sieve, the bases and densities extended to the sets

$$G_m^\sharp = \prod_{\ell|m} G_\ell^\sharp$$

for $m \in S(\Lambda)$ have a similar interpretation. Indeed, G_m^\sharp identifies clearly with the set of conjugacy classes of the finite group $G_m = \prod G_\ell$. The density v_m is therefore still given by

$$v_m(y^\sharp) = \frac{|y^\sharp|}{|G_m|}.$$

Also, it is well known that the irreducible representations of G_m are of the form

$$\pi : g \mapsto \boxtimes_{\ell|m} \pi_\ell(g)$$

for some uniquely defined irreducible representations π_ℓ of G_ℓ , where \boxtimes is the external tensor product defined by

$$g = (g_\ell) \mapsto \bigotimes_{\ell|m} \rho_\ell(g_\ell).$$

In other words, the set Π_m of irreducible linear representations of G_m is identified canonically with $\prod \Pi_\ell$. Moreover, the character of a representation of G_m of this form is simply

$$\mathrm{Tr} \pi(g) = \prod_{\ell|m} \mathrm{Tr} \pi_\ell(g_\ell),$$

so that the basis \mathcal{B}_m obtained from \mathcal{B}_ℓ is none other than the basis of functions $y^\sharp \mapsto \mathrm{Tr} \pi(y^\sharp)$ for π ranging over Π_m .

Given a siftable set (X, μ, F) associated to a conjugacy sieve $(G, \Lambda, (\rho_\ell))$, the sums $W(\varphi, \varphi')$ become

$$W(\pi, \tau) = \int_X \mathrm{Tr} \pi(\rho_m(F_x)) \overline{\mathrm{Tr} \tau(\rho_n(F_x))} d\mu(x) \quad (3.1)$$

for irreducible representations π and τ of G_m and G_n respectively, which can usually be interpreted as *exponential sums* (or integrals) over X , since the character values, as traces of matrices of finite order, are sums of finitely many roots of unity.

In Section 3.5, we include a self-contained statement of the conjugacy sieve for ease of reference.

3.2 Group sieves

The general sieve setting can also be applied to problems where the sieving sets are not conjugacy-invariant, using a basis of *matrix coefficients* of irreducible representations. Let $(G, \Lambda, (\rho_\ell))$ be a group sieve setting. For each ℓ and each irreducible representation $\pi \in \Pi_\ell$, choose an orthonormal basis $(e_{\pi,i})$ of the space V_π of the representation (with respect to a G_ℓ -invariant inner product $\langle \cdot, \cdot \rangle_\pi$). Then (see, e.g., [79, Section I.5], which treats compact groups), the family \mathcal{B}_ℓ of functions of the type

$$\begin{aligned} \varphi_{\pi,e,f} : x &\mapsto \sqrt{\dim \pi} \langle \pi(x)e, f \rangle_\pi, \\ e &= e_{\pi,1}, \dots, e_{\pi, \dim \pi}, \quad f = e_{\pi,1}, \dots, e_{\pi, \dim \pi} \end{aligned}$$

is an orthonormal basis of $L^2(G_\ell)$ for the inner product

$$\langle f, g \rangle = \frac{1}{|G_\ell|} \sum_{x \in G_\ell} f(x) \overline{g(x)},$$

i.e., corresponding to the uniform density $v_\ell(x) = 1/|G_\ell|$ for all $x \in G_\ell$. Moreover, for $\pi = 1$, and an arbitrary choice of $e \in \mathbf{C}$ with $|e| = 1$, the function $\varphi_{1,e,e} \in \mathcal{B}_\ell$ is the constant function 1.

If we extend the basis \mathcal{B}_ℓ to orthonormal bases \mathcal{B}_m of $L^2(G_m)$ for all $m \in S(\Lambda)$, by multiplicativity, the functions in \mathcal{B}_m are of the type

$$\varphi_{\pi,e,f} : (x_\ell) \mapsto \sqrt{\dim \pi} \prod_{\ell|m} \langle \pi_\ell(x_\ell) e_\ell, f_\ell \rangle_{\pi_\ell}$$

where $e = \otimes e_\ell$ and $f = \otimes f_\ell$ run over elements of the orthonormal basis

$$\left(\bigotimes_{\ell|m} e_{\pi_\ell, i_\ell} \right), \quad 1 \leq i_\ell \leq \dim \pi_\ell,$$

constructed from the chosen bases $(e_{\pi,i})$ of the components, the inner product on the space of $\boxtimes \pi_\ell$ being the natural G_m -invariant one:

$$\langle \otimes e_\ell, \otimes f_\ell \rangle = \prod_{\ell|m} \langle e_\ell, f_\ell \rangle.$$

The sums $W(\varphi, \varphi') = W(\varphi_{\pi,e,f}, \varphi_{\tau,e',f'})$ occurring in Proposition 2.9 to estimate the large sieve constant are given by

$$\sqrt{(\dim \pi)(\dim \tau)} \int_X \langle \pi(\rho_m(F_x))e, f \rangle_\pi \overline{\langle \tau(\rho_n(F_x))e', f' \rangle_\tau} d\mu(x). \quad (3.2)$$

If we apply Lemma 2.11 to elements $\varphi_{\pi,e,f}, \varphi_{\tau,e',f'}$ of the basis \mathcal{B}_m and \mathcal{B}_n of $L^2(G_m)$ and $L^2(G_n)$, the function $[\varphi_\pi, \overline{\varphi_\tau}]$ which is integrated can be written as a matrix coefficient of the representation

$$[\pi, \bar{\tau}] = \pi_{m'} \boxtimes (\pi_d \otimes \bar{\tau}_d) \boxtimes \overline{\tau_{n'}} \quad (3.3)$$

of $G_{[m,n]}$, where we write $\pi = \pi_{m'} \boxtimes \pi_d, \tau = \tau_{n'} \boxtimes \tau_d$, with the obvious meaning of the components $\pi_{m'}, \pi_d, \tau_d, \tau_{n'}$, and the bar indicates taking the contragredient representation.

Indeed, we have

$$[\varphi_{\pi,e,f}, \overline{\varphi_{\tau,e',f'}}](x_\ell) = \sqrt{(\dim \pi)(\dim \tau)} \langle [\pi, \bar{\tau}](\rho_{[m,n]}(F_x))\tilde{e}, \tilde{f} \rangle_{[\pi, \bar{\tau}]}$$

for $(x_\ell) \in G_{[m,n]}$, with $\tilde{e} = e \otimes e', \tilde{f} = f \otimes f'$.

Concretely, this means that in order to deal with the sums $W(\varphi, \varphi')$ to estimate the large sieve constant using the basis \mathcal{B}_m of matrix coefficients, it suffices to be able to estimate all integrals of the type

$$\int_X \langle \varpi(F_x)e, f \rangle_{\varpi} d\mu(x) \quad (3.4)$$

where ϖ is a representation of G that factors through a finite product of groups G_ℓ , and e, f are vectors in the space of the representation ϖ (the inner product

being G -invariant). See the proof of Part (2) of Theorem 7.4 for an application of this. Again, see Section 3.5 for a self-contained statement of the general group sieve.

Remark 3.1 Another potentially useful sieve setting associated to a group sieve setting (G, Λ, ρ_ℓ) is obtained by replacing ρ_ℓ with the projections $G \rightarrow G_\ell \rightarrow G_\ell/K_\ell = Y_\ell$ for $\ell \in \Lambda$, where K_ℓ is an arbitrary subgroup of G_ℓ . Considering the density on Y_ℓ which is the image of the uniform density on G_ℓ , an orthonormal basis \mathcal{B}_ℓ of $L^2(Y_\ell)$ is then obtained by taking the functions

$$\varphi_{\pi, e, f} : gK_\ell \mapsto \sqrt{\dim \pi} \langle \pi(g)e, f \rangle$$

where π runs over irreducible representations of G_ℓ , e runs over an orthonormal basis of the K_ℓ -invariant subspace in the space V_π of π , and f over a full orthonormal basis of V_π .

Indeed, the restriction on e ensures that such functions are well-defined on G_ℓ/K_ℓ (i.e., the matrix coefficient is K_ℓ -invariant), and since those are matrix coefficients, there only remains to check that they span $L^2(Y_\ell)$. However, using Frobenius reciprocity, the total number of such functions is

$$\sum_{\pi} (\dim \pi) \langle \text{Res}_{K_\ell}^{G_\ell} \pi, 1 \rangle_{K_\ell} = \sum_{\pi} (\dim \pi) \langle \pi, \text{Ind}_{K_\ell}^{G_\ell} 1 \rangle_{G_\ell} = \dim \text{Ind}_{K_\ell}^{G_\ell} 1 = |Y_\ell|$$

and since they are independent, the result follows.

Because this basis is a sub-basis of the previous one, any estimate for the large sieve constant for the group sieve will give one for this sieve setting.

3.3 Coset sieves

Our next subject, a generalization of conjugacy sieves, is the setting in which the sieve for Frobenius over finite fields of [80] and Chapter 8 operates.

We start again with a group G and a family of surjective homomorphisms $G \rightarrow G_\ell$ onto finite groups for $\ell \in \Lambda$. However, we also assume that there is a normal subgroup G^g of G such that the quotient G/G^g is abelian. Define $G_\ell^g = \rho_\ell(G^g)$, which is a normal subgroup of G_ℓ (since ρ_ℓ is surjective), and let $\Gamma_\ell = G_\ell/G_\ell^g$. We obtain a commutative diagram with exact rows

$$\begin{array}{ccccccc} 1 & \longrightarrow & G^g & \longrightarrow & G & \xrightarrow{d} & G/G^g \longrightarrow 1 \\ & & \downarrow & & \rho_\ell \downarrow & & p \downarrow \\ 1 & \longrightarrow & G_\ell^g & \longrightarrow & G_\ell & \xrightarrow{d} & \Gamma_\ell \longrightarrow 1, \end{array} \quad (3.5)$$

where the downward arrows are surjective. The groups Γ_ℓ are finite abelian groups.

We then extend this by multiplicativity as in the case of group sieves, putting

$$G_m = \prod_{\ell|m} G_\ell, \quad G_m^g = \prod_{\ell|m} G_\ell^g$$

for $m \in S(\Lambda)$. The group G_m^g is a normal subgroup of G_m , and we let $\Gamma_m = G_m/G_m^g$. Then we can still write commutative diagrams with exact rows

$$\begin{array}{ccccccc} 1 & \longrightarrow & G^g & \longrightarrow & G & \xrightarrow{d} & G/G^g \longrightarrow 1 \\ & & \downarrow & & \rho_m \downarrow & & p \downarrow \\ 1 & \longrightarrow & G_m^g & \longrightarrow & G_m & \xrightarrow{d} & \Gamma_m \longrightarrow 1, \end{array} \quad (3.6)$$

but the downward arrows are no longer necessarily surjective.

The sieve setting for a coset sieve is obtained by fixing some $\alpha \in G/G^g$, or equivalently some G^g -coset in G , and considering the set $Y \subset G^\sharp$ of G -conjugacy classes in the G^g -coset $d^{-1}(\alpha)$. Since G^g is normal in G , this coset $d^{-1}(\alpha)$ is indeed invariant under conjugation by the whole of G (this is an important point).

We then let $Y_\ell = \rho_\ell(Y) \subset G_\ell^\sharp$, and see that this is also the set of G_ℓ -conjugacy classes in the G_ℓ^g -coset defined by $p(\alpha) \in \Gamma_\ell$. Hence we have a sieve setting

$$(Y, \Lambda, (Y \xrightarrow{\rho_\ell} Y_\ell)).$$

The natural density to consider (which arises in the sieve for Frobenius) is still

$$v_\ell(y^\sharp) = \frac{|y^\sharp|}{|G_\ell^g|}, \quad \text{and hence} \quad v_m(y^\sharp) = \frac{|y^\sharp|}{|G_m^g|}$$

for a conjugacy class y^\sharp . Note that this means that for any conjugacy-invariant subset $\Omega_\ell \subset G_\ell$, union of a set Ω_ℓ^\sharp of conjugacy classes such that $\Omega_\ell^\sharp \subset d^{-1}(p(\alpha)) = Y_\ell$, we have

$$v(\Omega_\ell^\sharp) = \frac{|\Omega_\ell|}{|G_\ell^g|}.$$

We turn to the question of finding a suitable orthonormal basis of $L^2(Y_\ell, v_\ell)$. This is provided by the following general lemma, which applies equally to $H = G_\ell$ and to $H = G_m$ for $m \in S(\Lambda)$.

Lemma 3.2 *Let H be a finite group, H^g a normal subgroup with abelian quotient $\Gamma = H/H^g$. Let $\alpha \in \Gamma$ and let Y be the set of conjugacy classes of H with image α in Γ .*

For an irreducible linear representation π of H , let φ_π be the function

$$\varphi_\pi : y^\sharp \mapsto \text{Tr } \pi(y^\sharp)$$

on H^\sharp . Consider the inner product

$$\langle f, g \rangle = \frac{1}{|H^g|} \sum_{y^\sharp \in Y} |y^\sharp| f(y^\sharp) \overline{g(y^\sharp)}.$$

for functions f and g defined on Y .

(1) For π, τ irreducible linear representations of H , we have

$$\langle \varphi_\pi, \varphi_\tau \rangle = \begin{cases} 0, & \text{if either } \pi \upharpoonright H^g \not\simeq \tau \upharpoonright H^g \text{ or } \varphi_\pi \upharpoonright Y = 0, \\ \overline{\psi(\alpha)} |\hat{\Gamma}^\pi|, & \text{where } \psi \in \hat{\Gamma} \text{ satisfies } \pi \otimes \psi \simeq \tau, \text{ otherwise,} \end{cases} \quad (3.7)$$

where $\hat{\Gamma}$ is the group of characters of Γ and $\hat{\Gamma}^\pi = \{\psi \in \hat{\Gamma} \mid \pi \simeq \pi \otimes \psi\}$.

(2) Let Π be a set of representatives of irreducible linear representations of H for the equivalence relation

$$\pi \sim \tau \text{ if and only if } \pi \upharpoonright H^g \simeq \tau \upharpoonright H^g,$$

and let \mathcal{B} be the family of functions

$$\begin{cases} Y & \rightarrow \mathbf{C} \\ y^\sharp & \mapsto \frac{\varphi_\pi(y^\sharp)}{\sqrt{|\hat{\Gamma}^\pi|}}, \end{cases}$$

on Y , where π ranges over the subset $\Pi^* \subset \Pi$ where $\pi \in \Pi^*$ if and only if $\varphi_\pi \upharpoonright Y \neq 0$. Then \mathcal{B} is an orthonormal basis of $L^2(Y)$ for the above inner product.

In the second case of (3.7), the existence of the character ψ will follow from the proof below.

Proof We have

$$\begin{aligned} \langle \varphi_\pi, \varphi_\tau \rangle &= \frac{1}{|H^g|} \sum_{\substack{y \in H \\ d(y) = \alpha}} \text{Tr } \pi(y) \overline{\text{Tr } \tau(y)} \\ &= \frac{1}{|H^g|} \frac{1}{|\Gamma|} \sum_{y \in H} \left(\sum_{\psi \in \hat{\Gamma}} \overline{\psi(\alpha)} \psi(y) \right) \text{Tr } \pi(y) \overline{\text{Tr } \tau(y)} \\ &= \sum_{\psi \in \hat{\Gamma}} \overline{\psi(\alpha)} \langle \pi \otimes \psi, \tau \rangle_H = \sum_{\psi \in \hat{\Gamma}} \overline{\psi(\alpha)} \delta(\pi \otimes \psi, \tau), \end{aligned}$$

by orthogonality of characters of irreducible representations in $L^2(H)$.

First of all, this is certainly zero unless there exists at least one ψ such that $\pi \otimes \psi \simeq \tau$. In such a case we have $\pi|H^g \simeq \tau|H^g$ since $H^g \subset \text{Ker}(\psi)$, so we have shown that the condition $\pi|H^g \not\simeq \tau|H^g$ implies that the inner product is zero.

Assume now that $\pi|H^g \simeq \tau|H^g$; then repeating the above with $\alpha = 1$ (i.e., $Y = H^g$), it follows from $\langle \pi, \tau \rangle_{H^g} \neq 0$ that there exists one ψ at least such that $\pi \otimes \psi = \tau$.

Fixing one such character ψ_0 , the characters ψ' for which $\pi \otimes \psi' \simeq \tau$ are given by $\psi' = \psi \psi_0$ where $\psi \in \hat{\Gamma}^\pi$. Then we find

$$\langle \varphi_\pi, \varphi_\tau \rangle = \sum_{\psi \in \hat{\Gamma}^\pi} \overline{\psi(\alpha)} \delta(\pi \otimes \psi, \pi \otimes \psi_0) = \overline{\psi_0(\alpha)} \sum_{\psi \in \hat{\Gamma}^\pi} \overline{\psi(\alpha)}.$$

For any $\psi \in \hat{\Gamma}^\pi$ and $y^\sharp \in Y$, we have the character relation

$$\text{Tr } \pi(y^\sharp) = \psi(y^\sharp) \text{Tr } \pi(y^\sharp) = \psi(\alpha) \text{Tr } \pi(y^\sharp),$$

hence either $\psi(\alpha) = 1$ for all ψ , or $\text{Tr } \pi(y^\sharp) = 0$ for all y^\sharp , i.e., φ_π restricted to Y vanishes. In this last case, we have trivially $\varphi_\tau = 0$ also on Y , and the inner product is zero.

So we are led to the last case where $\pi|H^g = \tau|H^g$ but $\psi(\alpha) = 1$ for all $\psi \in \hat{\Gamma}^\pi$. Then the inner product formula is clear from the above.

Now to prove (2) from (1), notice first that the family \mathcal{B} is a generating set of $L^2(Y)$ (indeed, all φ_π generate $L^2(H^\sharp)$, but those π for which $\varphi_\pi = 0$ on Y are clearly not needed, and if $\pi \sim \tau$, we have $\varphi_\tau = \psi(\alpha)\varphi_\pi$ on Y , where ψ satisfies $\tau \simeq \psi \otimes \pi$, so one element of each equivalence class suffices for functions on Y). Then the fact that we have an orthonormal basis follows from the inner product formula, observing that if $\tau \simeq \pi \otimes \psi$, we have in fact $\pi = \tau$ by definition of the equivalence relation, so $\psi = 1$ in (3.7). \square

Example 3.3 In this lemma we emphasize that distinct representations of H may give the same restriction on H^g , in which case they correspond to a single element of the basis, and that it is possible that a φ_π vanish on Y , in which case the representative in question is discarded from the basis.

Take for instance $G = D_n$, a dihedral group of order $2n$. There is an exact sequence

$$1 \rightarrow \mathbf{Z}/n\mathbf{Z} \rightarrow G \xrightarrow{d} \mathbf{Z}/2\mathbf{Z} \rightarrow 1$$

and if $Y = d^{-1}(1) \subset G$ and π is any representation of G of degree 2, we have $\text{Tr } \pi(x) = 0$ for all $x \in Y$ (see, e.g. [115, 5.3]).

In particular, note that even though both cosets of $\mathbf{Z}/n\mathbf{Z}$ in G have n elements, the sets of conjugacy classes in each do not have the same cardinality (if n is

odd, there are $(n + 1)/2$ classes in $\text{Ker } d$ and 1 in the other coset, while if n is even, there are $n/2 + 1$ classes in $\text{Ker } d$ and 2 in the other coset). In other words, in a coset sieve, the spaces Y_m strongly depend on the value of α .

If we apply Lemma 3.2 to the groups G_m and their subgroups G_m^s , we clearly obtain orthonormal bases of $L^2(Y_m)$ containing the constant function 1, for the density ν_m above, and moreover, it is easily seen that they are obtained ‘multiplicatively’ from the case of G_ℓ . Although it was not phrased in this manner, this is what was used in [80].¹

We again include a self-contained statement in Section 3.5 (notice that in order to simplify matters a bit, we do not ask there for Π_m^* to exclude representations with character vanishing on Y_m , since they do not contribute to the left-hand side of the inequality defining the large sieve constant).

3.4 Exponential sums and equidistribution for group sieves

We now consider what happens with the equidistribution approach for coset sieves. (Hence also for conjugacy sieves, where $G^s = G$.)

If we apply Lemma 2.11 to the elements $\varphi_\pi, \varphi_\tau$ of the bases \mathcal{B}_m and \mathcal{B}_n of $L^2(Y_m)$, we see that the function $[\varphi_\pi, \overline{\varphi_\tau}]$ defined in (2.9) is the character of the representation

$$[\pi, \bar{\tau}] = \pi_{m'} \boxtimes (\pi_d \otimes \bar{\tau}_d) \boxtimes \bar{\tau}_{n'} \quad (3.8)$$

of $G_{[m,n]}$, already defined in (3.3). Hence we have

$$W(\pi, \tau) = \frac{1}{\sqrt{|\hat{\Gamma}_m^\pi| |\hat{\Gamma}_n^\tau|}} \int_X \text{Tr}([\pi, \bar{\tau}] \rho_{[m,n]}(F_x)) d\mu(x). \quad (3.9)$$

In applications, this means that to estimate the integrals $W(\pi, \tau)$ it suffices (and may be more convenient) to be able to deal with integrals of the form

$$\int_X \text{Tr}(\varpi(F_x)) d\mu(x)$$

where ϖ is a representation of G that factors through a finite product of groups G_ℓ (see Chapter 7 for an instance of this).

¹ With minor differences, e.g., the upper bound κ for the order of $\hat{\Gamma}_m^\pi$ that occurs in [80], and can be removed – as also noticed independently by Zywinia in a private email.

Note that if we approach these integrals using the equidistribution method, then the analogue of (2.10) is the identity

$$|\{\rho_d(F_x) = y^\sharp\}| = \int_{\{\rho_d(F_x) = y^\sharp\}} d\mu(x) = \frac{|y^\sharp|}{|G_d^g|} |X| + r_d(X; y^\sharp), \quad (3.10)$$

defining $r_d(X; y^\sharp)$ for $y^\sharp \in Y_d$. Then (2.11) becomes

$$W(\pi, \tau) = \frac{|X|}{\sqrt{|\hat{\Gamma}_m^\pi| |\hat{\Gamma}_n^\tau|}} m([\pi, \bar{\tau}]) + O\left(\frac{1}{\sqrt{|\hat{\Gamma}_m^\pi| |\hat{\Gamma}_n^\tau|}} \sum_{y^\sharp \in Y_{[m,n]}} \dim[\pi, \bar{\tau}] |r_{[m,n]}(X; y^\sharp)|\right)$$

(using the trivial bound

$$|\mathrm{Tr} \pi(x)| \leq \dim \pi$$

for the absolute value of the character of a representation π). By Lemma 2.12, we have $m([\pi, \bar{\tau}]) = \delta((m, \pi), (n, \tau))$ (see also Lemma 3.2 with $H = G_{[m,n]}$).

Using this and (2.11), we get

$$W(\pi, \tau) = \delta(\pi, \tau) |X| + O\left(\sum_{y^\sharp \in Y_{[m,n]}} \dim[\pi, \bar{\tau}] |r_{[m,n]}(X; y^\sharp)|\right),$$

where the implied constant is ≤ 1 . Hence for any sieve support \mathcal{L} , the large sieve bound of Proposition 2.9 holds with

$$\Delta \leq |X| + R(X; \mathcal{L}) \quad (3.11)$$

where

$$R(X; \mathcal{L}) = \max_{n \in \mathcal{L}} \max_{\pi \in \Pi_n^*} \left\{ \sum_{m \in \mathcal{L}} \sum_{\tau \in \Pi_m^*} \sum_{y^\sharp \in Y_{[m,n]}} \dim[\pi, \bar{\tau}] |r_{[m,n]}(X; y^\sharp)| \right\}. \quad (3.12)$$

For later reference, we also note the following fact:

Lemma 3.4 *Let m, n be in $S(\Lambda)$, $\pi \in \Pi_m^*$, $\tau \in \Pi_n^*$. The multiplicity of the trivial representation in the restriction of $[\pi, \bar{\tau}]$ to $G_{[m,n]}^g$ is equal to zero if $(m, \pi) \neq (n, \tau)$, and is equal to $|\hat{\Gamma}_m^\pi|$ if $(m, \pi) = (n, \tau)$.*

Proof This multiplicity is by definition $\langle [\pi, \bar{\tau}], 1 \rangle$ computed in $L^2(G_{[m,n]}^g)$, i.e., it is $\langle \varphi_\pi, \varphi_\tau \rangle$ in $L^2(Y_{[m,n]})$ in the case $\alpha = 1 \in G/G^g$ (where π and τ , and hence φ_π and φ_τ , are extended to $G_{[m,n]}$ by adding trivial components at $\ell \notin m$ or $\ell \notin n$, respectively). So the result is a consequence of Lemma 3.2. \square

3.5 Self-contained statements

The sieves described in this chapter are likely to be the most commonly used in applications. In order to ease further references, we conclude with self-contained statements which do not involve any new terminology.

Proposition 3.5 *Let G be a group, Λ a set, and for $\ell \in \Lambda$, let $\rho_\ell : G \rightarrow G_\ell$ be a surjective map onto a finite group. Moreover, let (X, μ) be a finite measure space and $F : X \rightarrow G$ a map such that $\{x | \rho_\ell(F_x) = y\}$ is measurable for all $\ell \in \Lambda$ and $y \in G_\ell$.*

For $m \subset \Lambda$, let

$$G_m = \prod_{\ell \in m} G_\ell,$$

and let Π_m^ be the set of primitive irreducible linear representations of G_m , i.e., those such that no component π_ℓ for $\ell \in m$ is trivial when writing*

$$\pi \simeq \boxtimes_{\ell \in m} \pi_\ell.$$

Let \mathcal{L}^ be a finite subset of Λ , \mathcal{L} a finite collection of subsets of \mathcal{L}^* . Then, for any conjugacy invariant subsets $\Omega_\ell \subset G_\ell$ for $\ell \in \mathcal{L}^*$, we have*

$$\mu(\{x \in X | \rho_\ell(F_x) \notin \Omega_\ell, \text{ for } \ell \in \mathcal{L}^*\}) \leq \Delta H^{-1}$$

where Δ is the smallest non-negative real number such that

$$\sum_{m \in \mathcal{L}} \sum_{\pi \in \Pi_m^*} \left| \int_X \alpha(x) \operatorname{Tr} \pi(\rho_m(F_x)) d\mu(x) \right|^2 \leq \Delta \int_X |\alpha(x)|^2 d\mu(x)$$

for all square-integrable functions $\alpha \in L^2(X, \mu)$, and

$$H = \sum_{m \in \mathcal{L}} \prod_{\ell \in m} \frac{|\Omega_\ell|}{|G_\ell| - |\Omega_\ell|}.$$

Moreover we have

$$\Delta \leq \max_{m \in \mathcal{L}} \max_{\pi \in \Pi_m^*} \sum_{n \in \mathcal{L}} \sum_{\tau \in \Pi_n^*} |W(\pi, \tau)|,$$

where

$$W(\pi, \tau) = \int_X \operatorname{Tr} \pi(\rho_m(F_x)) \overline{\operatorname{Tr} \tau(\rho_n(F_x))} d\mu(x).$$

Proposition 3.6 *Let $(G, \Lambda, (\rho_\ell))$, (X, μ, F) be as in Proposition 3.5, and define G_m and Π_m^* as above. Moreover, for each $\ell \in \Lambda$ and each $\pi \in \Pi_\ell^*$, let*

$$(e_{\pi,1}, \dots, e_{\pi, \dim \pi})$$

be an orthonormal basis of the space of π with respect to a G_ℓ -invariant inner product, and for any finite subset $m \subset \Lambda$ and $\pi \in \Pi_m^*$, fix an isomorphism

$$\pi \simeq \boxtimes_{\ell \in m} \pi_\ell,$$

and let

$$(e_{\pi,1}, \dots, e_{\pi, \dim \pi})$$

denote the orthonormal basis of the space of π obtained by tensor product of those of the components.

Let \mathcal{L}^* be a finite subset of Λ , \mathcal{L} a finite collection of subsets of \mathcal{L}^* . Then, for any subsets $\Omega_\ell \subset G_\ell$ for $\ell \in \mathcal{L}^*$, we have

$$\mu(\{x \in X \mid \rho_\ell(F_x) \notin \Omega_\ell, \text{ for } \ell \in \mathcal{L}^*\}) \leq \Delta H^{-1}$$

where Δ is the smallest non-negative real number such that

$$\begin{aligned} & \sum_{m \in \mathcal{L}} \sum_{\pi \in \Pi_m^*} \sqrt{\dim(\pi)} \sum_{i,j=1}^{\dim \pi} \left| \int_X \alpha(x) \langle \pi(\rho_m(F_x)) e_{\pi,i}, e_{\pi,j} \rangle d\mu(x) \right|^2 \\ & \leq \Delta \int_X |\alpha(x)|^2 d\mu(x) \end{aligned}$$

for all square-integrable functions $\alpha \in L^2(X, \mu)$, where

$$H = \sum_{m \in \mathcal{L}} \prod_{\ell \in m} \frac{|\Omega_\ell|}{|G_\ell| - |\Omega_\ell|}.$$

Moreover we have

$$\Delta \leq \max_{m \in \mathcal{L}} \max_{\pi \in \Pi_m^*} \max_{i,j \leq \dim \pi} \sum_{n \in \mathcal{L}} \sum_{\tau \in \Pi_n^*} \sum_{k,l \leq \dim \tau} \sqrt{\dim(\pi) \dim(\tau)} |W(\pi, i, j; \tau, k, l)|,$$

where

$$W(\pi, i, j; \tau, k, l) = \int_X \langle \pi(\rho_m(F_x)) e_{\pi,i}, e_{\pi,j} \rangle \overline{\langle \tau(\rho_n(F_x)) e_{\tau,k}, e_{\tau,l} \rangle} d\mu(x).$$

Proposition 3.7 *Let G be a group, G^g a normal subgroup of G with abelian quotient G/G^g ; denote by $d: G \rightarrow G/G^g$ the quotient map. Let Λ be a group and let $\rho_\ell: G \rightarrow G_\ell$, for $\ell \in \Lambda$, be a family of surjective homomorphisms onto finite groups. Denote $G_\ell^g = \rho_\ell(G^g)$. Let $\alpha \in G/G^g$ be fixed, and let $Y = d^{-1}(\alpha) \subset G$. Let (X, μ) be a finite measure space and $F: X \rightarrow Y$ a map such that $\{x \mid \rho_\ell(F_x) = y\}$ is measurable for all $\ell \in \Lambda$ and $y \in G_\ell$. For any subset $m \in \Lambda$, let*

$$G_m = \prod_{\ell \in m} G_\ell, \quad G_m^g = \prod_{\ell \in m} G_\ell^g,$$

and let Π_m be a set of representatives of the set of irreducible representations of G_m modulo equality restricted to G_m^s , containing the constant function 1. Moreover, let Π_m^* be the subset of primitive representations, i.e., those such that when π is decomposed as $\boxtimes_{\ell \in m} \pi_\ell$, no component π_ℓ is trivial. Let $\hat{\Gamma}_m^\pi$ be the set of characters ψ of G_m/G_m^s such that $\pi \otimes \psi \simeq \pi$ for a representation π of G_m .

Let \mathcal{L}^* be a finite subset of Λ , \mathcal{L} a finite collection of subsets of \mathcal{L}^* . Then, for any conjugacy invariant subsets $\Omega_\ell \subset G_\ell$ for $\ell \in \mathcal{L}^*$, we have

$$\mu(\{x \in X \mid \rho_\ell(F_x) \notin \Omega_\ell, \text{ for } \ell \in \mathcal{L}^*\}) \leq \Delta H^{-1}$$

where Δ is the smallest non-negative real number such that

$$\sum_{m \in \mathcal{L}} \sum_{\pi \in \Pi_m^*} \left| \int_X \alpha(x) \operatorname{Tr} \pi(\rho_m(F_x)) d\mu(x) \right|^2 \leq \Delta \int_X |\alpha(x)|^2 d\mu(x)$$

for all square-integrable functions $\alpha \in L^2(X, \mu)$, and

$$H = \sum_{m \in \mathcal{L}} \prod_{\ell \in m} \frac{|\Omega_\ell|}{|G_\ell^s| - |\Omega_\ell|}.$$

In addition, we have

$$\Delta \leq \max_{m \in \mathcal{L}} \max_{\pi \in \Pi_m^*} \sum_{n \in \mathcal{L}} \sum_{\tau \in \Pi_n^*} |W(\pi, \tau)|, \quad (3.13)$$

with

$$W(\pi, \tau) = \frac{1}{\sqrt{|\hat{\Gamma}_m^\pi| |\hat{\Gamma}_n^\tau|}} \int_X \operatorname{Tr} \pi(\rho_m(F_x)) \overline{\operatorname{Tr} \tau(\rho_n(F_x))} d\mu(x). \quad (3.14)$$

4

Elementary and classical examples

This intermediate chapter describes how classical forms of the large sieve are special cases of the setting described in Chapter 2, starting with the enlightening (if not particularly useful) example of the inclusion-exclusion principle which is also often used in probability theory and combinatorics. We have made no special effort to be exhaustive, and in particular we do not try to survey the early applications of the large sieve, many of which can be found by browsing through issues of the journal *Mathematika* from the 1950s and 1960s.

4.1 The inclusion-exclusion principle

The first example illustrates the general sieve setting, showing that it includes (and extends) the inclusion-exclusion familiar in combinatorics and probability theory, and also that the large sieve inequality is sharp in this general context (i.e., there may be equality $|S(X, \Omega; \mathcal{L}^*)| = \Delta H^{-1}$).

Let $(\Omega, \Sigma, \mathbf{P})$ be a probability space and $A_\ell \subset \Sigma$, for $\ell \in \Lambda$, a countable family of events. Consider the event

$$A = \{\omega \in \Omega \mid \omega \notin A_\ell \text{ for any } \ell \in \Lambda\}.$$

For $m \in S(\Lambda)$, denote

$$A_m = \bigcap_{\ell \in m} A_\ell, \quad A_\emptyset = \Omega.$$

If Λ is finite, which we now assume, the inclusion-exclusion formula is

$$\mathbf{P}(A) = \sum_{m \in S(\Lambda)} (-1)^{|m|} \mathbf{P}(A_m),$$

and in particular, if the events are independent (as a whole), we have

$$\mathbf{P}(A_m) = \prod_{\ell \in m} \mathbf{P}(A_\ell), \quad \text{and} \quad \mathbf{P}(A) = \prod_{\ell \in \Lambda} (1 - \mathbf{P}(A_\ell)).$$

Take the sieve setting $(\Omega, \Lambda, \mathbf{1}_{A_\ell})$, where $\mathbf{1}_B$ is the characteristic function of an event B , with $Y_\ell = \{0, 1\}$ for all ℓ , and the siftable set $(\Omega, \mathbf{P}, \text{Id})$. Choose the density $v_\ell = \mathbf{1}_{A_\ell}(P)$, i.e., put

$$v_\ell(1) = \mathbf{P}(A_\ell), \quad v_\ell(0) = 1 - \mathbf{P}(A_\ell).$$

With sieving sets $\Omega_\ell = \{1\}$ for $\ell \in \Lambda$, we have precisely $S(X, \Omega; \Lambda) = A$.

The large sieve inequality yields

$$\mathbf{P}(A) \leq \Delta H^{-1}$$

where

$$H = \sum_{m \in \mathcal{L}} \prod_{\ell \in m} \frac{\mathbf{P}(A_\ell)}{1 - \mathbf{P}(A_\ell)},$$

and Δ is the large sieve constant for the sieve support \mathcal{L} , which may be any collection of subsets of Λ .

Coming to the large sieve constant, note that $L_0^2(Y_\ell)$ is one-dimensional for all ℓ , hence so is $L_0^2(Y_m)$ for all m (including $m = \emptyset$). The basis function φ_ℓ for $L_0^2(Y_\ell)$ (up to multiplication by a complex number with modulus 1) is given by

$$\varphi_\ell(y) = \frac{y - p_\ell}{\sqrt{p_\ell(1 - p_\ell)}}$$

where $p_\ell = \mathbf{P}(A_\ell)$ for simplicity, so that

$$\varphi_\ell(\mathbf{1}_{A_\ell}) = \frac{\mathbf{1}_{A_\ell} - \mathbf{P}(A_\ell)}{\sqrt{\mathbf{V}(\mathbf{1}_{A_\ell})}},$$

and in particular

$$\mathbf{E}(\varphi_\ell(\mathbf{1}_{A_\ell})) = \langle \varphi_\ell, 1 \rangle = 0, \quad \mathbf{E}(\varphi_\ell(\mathbf{1}_{A_\ell})^2) = \|\varphi_\ell\|^2 = 1.$$

Hence, for $\ell, \ell' \in \Lambda$, $W(\varphi_\ell, \varphi_{\ell'})$ is given by

$$W(\varphi_\ell, \varphi_{\ell'}) = \mathbf{E}(\varphi_\ell(\mathbf{1}_{A_\ell})\varphi_{\ell'}(\mathbf{1}_{A_{\ell'}})),$$

and it is (by definition) the correlation coefficient of the random variables $\mathbf{1}_{A_\ell}$ and $\mathbf{1}_{A_{\ell'}}$; explicitly

$$W(\varphi_\ell, \varphi_{\ell'}) = \begin{cases} 1 & \text{if } \ell = \ell', \\ \frac{\mathbf{P}(A_\ell \cap A_{\ell'}) - \mathbf{P}(A_\ell)\mathbf{P}(A_{\ell'})}{\sqrt{p_\ell(1 - p_\ell)p_{\ell'}(1 - p_{\ell'})}} & \text{otherwise.} \end{cases}$$

If (and only if) the (A_ℓ) form a family of pairwise independent events, we see that $W(\varphi_\ell, \varphi_{\ell'}) = \delta(\ell, \ell')$. More generally, in all cases, for any $m, n \subset \Lambda$, we have

$$W(\varphi_m, \varphi_n) = \mathbf{E} \left(\prod_{\ell \in m} \varphi_\ell(\mathbf{1}_{A_\ell}) \prod_{\ell \in n} \varphi_\ell(\mathbf{1}_{A_\ell}) \right),$$

which is a multiple normalized centred moment of the $\mathbf{1}_{A_\ell}$.

If the (A_ℓ) are globally independent, we obtain

$$\begin{aligned} W(\varphi_m, \varphi_n) &= \prod_{\substack{\ell \in m \cup n \\ \ell \notin m \cap n}} \frac{\mathbf{E}(\mathbf{1}_{A_\ell} - p_\ell)}{\sqrt{\mathbf{V}(\mathbf{1}_{A_\ell})}} \prod_{\ell \in m \cap n} \frac{\mathbf{E}((\mathbf{1}_{A_\ell} - p_\ell)^2)}{\sqrt{\mathbf{V}(\mathbf{1}_{A_\ell})}} \\ &= \delta(m, n) \end{aligned}$$

(since the first factor vanishes if the product is not empty, i.e., if $m \neq n$, and the second term is 1 by orthonormality of φ_ℓ). It follows by (2.8) that $\Delta \leq 1$, and in fact there must be equality. Moreover, in this situation, if \mathcal{L} contains all subsets of Λ , we have

$$H = \prod_{\ell \in \Lambda} \left(1 + \frac{p_\ell}{1 - p_\ell} \right) = \prod_{\ell \in \Lambda} \frac{1}{1 - p_\ell},$$

so that we find

$$\Delta H^{-1} \leq \prod_{\ell \in \Lambda} (1 - \mathbf{P}(A_\ell)) = \mathbf{P}(A),$$

i.e., the large sieve inequality is an equality here.

Similarly, the inequality (2.13) becomes an equality if the events are pairwise independent, and reflects the formula for the variance of a sum of (pairwise) independent random variables.

In the general case of possibly dependent events, on the other hand, we have a quantitative inequality for $\mathbf{P}(A)$ which may be of some interest (and may be already known!). In fact, we have several possibilities depending on the choice of sieve support. It would be interesting to determine if those inequalities are of some use in probability theory.

To conclude this example, note that *any* sieve, once the prime sieve support \mathcal{L}^* and the sieving sets (Ω_ℓ) are chosen, may be considered as a similar ‘binary’ sieve with $Y_\ell = \{0, 1\}$ (or $\{0\}$ if $\Omega_\ell = \emptyset$, or $\{1\}$ if $\Omega_\ell = Y_\ell$) for all ℓ , by replacing the sieve setting $(Y, \Lambda, (\rho_\ell))$ with $(Y, \mathcal{L}^*, \mathbf{1}_{\Omega_\ell})$.

4.2 The classical large sieve

We have already mentioned during the course of Chapter 2 that the classical large sieve arises from the (group) sieve setting

$$\Psi = (\mathbf{Z}, \{\text{primes}\}, \mathbf{Z} \rightarrow \mathbf{Z}/\ell\mathbf{Z})$$

where the condition for an additive character $x \mapsto e(ax/m)$ of $G_m = (\mathbf{Z}/m\mathbf{Z})$ to be primitive is equivalent with the classical condition that $(a, m) = 1$.

In the most typical case, the siftable sets are

$$X = \{n \geq 1 \mid N \leq n < N + M\}$$

with $F_x = x$, and the abstract sieving problem becomes the ‘original’ one of finding integers in X which lie outside certain residue classes modulo some primes ℓ .

More generally, take

$$\Psi = (\mathbf{Z}', \{\text{primes}\}, \mathbf{Z}' \rightarrow (\mathbf{Z}/\ell\mathbf{Z})^r)$$

(the reduction maps) and $X = \{(a_1, \dots, a_r) \mid N_i \leq a_i < N_i + M_i\}$, with F the identity map again. Then what results is the higher-dimensional large sieve (see, e.g., [46], [63]).

For completeness, we recall the estimates available for the large sieve constant in those two situations, when we take \mathcal{L}^* to be the set of primes $\leq L$, and \mathcal{L} to be the set of squarefree integers $\leq L$, for some $L \geq 1$. We write $S(X, \Omega; L)$ instead of $S(X, \Omega; \mathcal{L}^*)$.

Theorem 4.1 *With notation as above, we have $\Delta \leq N - 1 + L^2$ for $r = 1$ and $\Delta \leq \prod (\sqrt{N_i} + L)^2$ for all $r \geq 1$. In fact, we have*

$$\sum_{m \leq L} \sum_{\substack{a \in \mathbf{Z}/m\mathbf{Z} \\ (a, m) = 1}} \left| \sum_{M < n \leq M+N} a_n e\left(\frac{an}{m}\right) \right|^2 \leq (N - 1 + L^2) \sum_{M < n \leq M+N} |a_n|^2,$$

$$\sum_{m \leq L} \sum_{\substack{(a_i) \in (\mathbf{Z}/m\mathbf{Z})^r \\ (a_i, m) = 1}} \left| \sum_{M_i < n_i \leq M_i + N_i} b_n e\left(\frac{\langle a, n \rangle}{m}\right) \right|^2 \leq \prod_{i=1}^r (\sqrt{N_i} + L)^2 \sum_{M_i < n_i \leq M_i + N_i} |a_n|^2$$

for arbitrary complex numbers (a_n) , respectively complex vectors (b_n) with $b_n \in \mathbf{C}^r$. Note that the sum over m is not restricted to squarefree numbers.

In particular, for any sieve problem associated to the sieve setting above, we have

$$|S(X, \Omega; L)| \leq (N - 1 + L^2)H^{-1}, \quad \text{if } r = 1,$$

$$|S(X, \Omega; L)| \leq \prod_{i=1}^r (\sqrt{N_i} + L)^2 H^{-1}, \quad \text{if } r \geq 1,$$

where¹

$$H = \sum_{m \leq L}^b \prod_{\ell | m} \frac{|\Omega_\ell|}{\ell^r - |\Omega_\ell|}.$$

Proof for $r = 1$, in a weaker form For completeness, we provide a proof of a weaker form of the large sieve inequality in the one-variable case, namely

$$\Delta \leq 2\pi N + Q^2,$$

following the nice trick of Gallagher. In almost all applications, this is as strong as the inequality as stated above.

We assume $M = 0$: the general case is deduced by an obvious shift. The first step is the inequality

$$|f(\tfrac{1}{2})| \leq \int_0^1 (|f(t)| + \tfrac{1}{2}|f'(t)|)dt$$

for any smooth function f on $[0, 1]$, which is a consequence of the simple formula

$$f(\tfrac{1}{2}) = \int_0^1 f(t)dt + \int_0^{1/2} tf'(t)dt + \int_{1/2}^1 (t-1)f'(t)dt.$$

By a change of variable, this leads to

$$|f(x)| \leq \frac{1}{\delta} \int_{x-\delta/2}^{x+\delta/2} |f(t)|dt + \frac{1}{2} \int_{x-\delta/2}^{x+\delta/2} |f'(t)|dt$$

for f smooth on $[x - \delta/2, x + \delta/2]$.

This is applied to

$$f(t) = \left(\sum_{1 \leq n \leq N} a_n e(nt) \right)^2,$$

¹ Recall that the notation \sum^b indicates a sum restricted to squarefree numbers.

with $x = a/q$, $q \leq Q$ and $(a, q) = 1$, taking $\delta = Q^{-2}$. One finds

$$\begin{aligned} \left| \sum_{1 \leq n \leq N} a_n e\left(\frac{an}{q}\right) \right|^2 &\leq \int_{a/q-\delta/2}^{a/q+\delta/2} \left| \sum_{1 \leq n \leq N} a_n e(nt) \right|^2 dt \\ &\quad + \int_{a/q-\delta/2}^{a/q+\delta/2} \left| \left(\sum_{1 \leq n \leq N} a_n e(nt) \right) \left(2i\pi \sum_{1 \leq n \leq N} na_n e(nt) \right) \right| dt. \end{aligned}$$

Now the point is that the intervals $]a/q - \delta/2, a/q + \delta/2[$, for $q \leq Q$ (not necessarily squarefree!) and a coprime with q , are all disjoint. Summing over q and a , and using periodicity and positivity, this leads to

$$\begin{aligned} \sum_{q \leq Q} \sum_{(a,q)=1} \left| \sum_{1 \leq n \leq N} a_n e\left(\frac{an}{q}\right) \right|^2 &\leq \delta^{-1} \int_0^1 \left| \sum_{1 \leq n \leq N} a_n e(nt) \right|^2 dt \\ &\quad + \int_0^1 \left| \left(\sum_{1 \leq n \leq N} a_n e(nt) \right) \left(2i\pi \sum_{1 \leq n \leq N} na_n e(nt) \right) \right| dt, \end{aligned}$$

and applying Parseval's identity and then the Cauchy–Schwarz inequality, this is bounded by

$$Q^2 \sum_n |a_n|^2 + 2\pi \left(\sum_n |a_n|^2 \right)^{1/2} \left(\sum_n n^2 |a_n|^2 \right)^{1/2} \leq (Q^2 + 2\pi N) \sum_n |a_n|^2.$$

□

In the one-variable case, the first version with this formulation (also the first version of the large sieve with comparable strength) is due to Bombieri, though Roth had an earlier form of the ‘dual sieve’ with almost the same quality. The bound $N - 1 + L^2$ for the large sieve constant is due to Selberg (see, e.g., [67, Section 7.5] for a proof). Although this is not our main concern, it should be mentioned that there are many subtleties behind this classical inequality; see for instance O. Ramaré’s investigations [106] of the distribution of the eigenvalues of the underlying finite-dimensional operator (see Remark 2.5).

The higher-dimensional case as stated is due to Huxley, see [63]; other versions exist, but they may not be as efficient when it comes to the dependency on r , which we will have cause to exploit.

Note that the usual modern treatments of the large sieve deduce such estimates from an analytic inequality which is more general than the ones we used, namely

(for $r = 1$), the inequality

$$\sum_r \left| \sum_{M < n \leq M+N} a_n e(n\xi_r) \right|^2 \leq (N - 1 + \delta^{-1}) \sum_n |a_n|^2 \quad (4.1)$$

for arbitrary sets (ξ_r) of elements in \mathbf{R}/\mathbf{Z} which are δ -spaced, i.e., the distance $d(\xi_r, \xi_s)$ in \mathbf{R}/\mathbf{Z} is at least δ if $r \neq s$ (this was first considered by Davenport and Halberstam [26]; see also [67, Theorem 7.7], [98]). Then one observes, as is clear in the argument above, that the points a/q for $q \leq Q$ (squarefree or not) and $(a, q) = 1$, are δ -spaced with $\delta = Q^{-2}$.

The inequality (4.1) amounts to the consideration of the sums

$$\sum_{M < n \leq M+N} e((\xi_r - \xi_s)n) = W(\pi_r, \pi_s)$$

where $\pi_r : n \mapsto e(n\xi_r)$ and π_s are representations of $G = \mathbf{Z}$ which *do not* factor through a finite index subgroup. In particular, such sums can not be approached by using equidistribution in the image group. However, this also suggests trying to prove similar inequalities for general groups sieves, i.e., essentially, consider integrals (3.4) for arbitrary (unitary) representations ϖ of G .

Note that for $r = 1$, the equidistribution assumption (2.10) becomes

$$\sum_{\substack{N \leq n < N+M \\ n \equiv y \pmod{d}}} 1 = \frac{M}{d} + r_d(X; y),$$

which holds with $|r_d(X; y)| \leq 1$ for any $y \in \mathbf{Z}/d\mathbf{Z}$. From (3.12) we obtain the estimate $\Delta \leq N + L^4$, which is by no means ridiculous.

Classical sieve theory is founded on such assumptions as (2.10), usually stated merely for $y = 0$, and on further assumptions concerning the resulting level of distribution, i.e., bounds for $r_d(X; 0)$ on average over d in a range as large as possible (compared with the size of X). More general bounds for $r_d(X; y)$ do occur however.

Note that, even if this is classical, the general framework clearly shows that to sieve an arbitrary set of integers $X \subset \{n \mid n \geq 1\} \subset \mathbf{Z}$, it suffices (at least up to a point!) to have estimates for exponential sums

$$\sum_{x \in X} e\left(\frac{ax}{m} - \frac{bx}{n}\right)$$

with n, m squarefree and $(a, m) = (b, n) = 1$. It suffices, in particular, to have equidistribution of X in (all) arithmetic progressions. This means for instance that some measure of large sieve is usually feasible for any sequence for which

the classical ‘small’ sieves work. This is of particular interest if X is sparse, in the sense that, e.g., $X \subset \{n \mid N < n \leq 2N\}$ for some N with $|X|/N$ going to zero.

It would also be interesting, as a problem in itself, to investigate the values of the large sieve constant when using other sieve support than squarefree integers up to L , for instance when the sieve support is the support of a combinatorial sieve (see, e.g., [67, 6.2]).

There are many applications of the classical form of the large sieve inequality. We highlight here a result of Gallagher [46] on the ‘generic’ irreducibility and maximality of the Galois group for integral polynomials with bounded height, because it is a reference point and motivation for some of the results obtained with more delicate sieve settings in Chapters 7 and 8 (and in [80]). Also, it is a very good example of a fairly direct application of the large sieve inequality.

Theorem 4.2 (Gallagher) *Let $r \geq 1$ be an integer. For any integer $N \geq 1$, let $E_r(N)$ be the set of monic polynomials of degree r in $\mathbf{Z}[T]$ such that*

$$f = T^r + a_{r-1}T^{r-1} + \cdots + a_1T + a_0$$

with $|a_i| \leq N$ for all i , which have the property that the splitting field K_f of f , i.e., the finite Galois extension of \mathbf{Q} generated by all roots of f , has Galois group G strictly smaller than the symmetric group \mathfrak{S}_r . Then we have

$$|E_r(N)| \ll r^3(2N+1)^{r-1/2}(\log N)$$

for $N \geq 2$, where the implied constant is absolute.

To be precise, the uniform result (with respect to the degree r) is a refinement of Gallagher’s result (valid for fixed r), and is stated in [80, Remark 7.4], with a small mistake (r^3 is replaced by r^2). Note that $(2N+1)^r$ is of course the exact number of monic polynomials in $\mathbf{Z}[T]$ with degree r and coefficients bounded by N . As explained by Gallagher in the introduction to his paper, the first results along those lines are due to Dörge and van der Waerden.

It is expected that the correct order of magnitude for $E_r(N)$ is N^{r-1} , up to logarithmic or similar factors, but no improvement on the exponent $r - 1/2$ has been obtained since Gallagher. Such a result would be an important breakthrough in the study of the large sieve.

Proof We give the details of the proof, as far as bounding the number of reducible polynomials (a weaker statement, where the factor r^3 is replaced by r^2), and sketch the extra ingredients required to control the splitting field. This will enable us to introduce some local counting results for polynomials over finite

fields which will be used again when dealing with sieves involving characteristic polynomials of matrices (see Chapters 7 and 8, as well as Appendix B).

We can obviously assume $r \geq 2$. The sieve setting we use is the r -dimensional one of Theorem 4.1, and the siftable set $X = X_r(N)$ is of course the set of all monic polynomials of degree r in $\mathbf{Z}[T]$ with coefficients $|a_i| \leq N$ for all i , $0 \leq i \leq r-1$. Now the crucial fact with respect to irreducibility is the following observation: if a polynomial $f \in X_r(N)$ is reducible (in $\mathbf{Q}[T]$, or equivalently in $\mathbf{Z}[T]$), then for any prime ℓ , the reduction of f modulo ℓ can not lie in the set Ω_ℓ corresponding to coefficients $(a_i) \in (\mathbf{F}_\ell)^r$ of an *irreducible* polynomial. So, the number of reducible polynomials is bounded by the size of the sifted set $S(X_r(N), \Omega; \mathcal{L}^*)$ for arbitrary finite set of primes \mathcal{L}^* .

Selecting $\mathcal{L}^* = \{\ell \leq L\}$ for some $L \geq 2$, and $\mathcal{L} = \mathcal{L}^*$ (with the usual identification), the large sieve inequality implies

$$|\{f \in X_r(N) \mid f \text{ is reducible}\}| \leq (\sqrt{2N+1} + L)^{2r} H^{-1}$$

with

$$H \geq \sum_{\ell \leq L} \frac{|\Omega_\ell|}{\ell^r}.$$

By Lemma B.1 of Appendix B, we have

$$H \geq \frac{1}{r} \sum_{4r < \ell \leq L} \left(1 - \frac{1}{\ell}\right)$$

if $L > 4r$, hence

$$H \geq \frac{\pi(L)}{r} + O(\log \log L) + O(1) \gg \frac{1}{r} \frac{L}{\log L}$$

for $L \geq \alpha r(\log 2r)$, where α is an absolute constant,² and the implied constant depends only on α , hence is absolute too.

Assuming that $\sqrt{2N+1} \geq \alpha r^2(\log 2r)$, we select $L = r^{-1}\sqrt{2N+1}$, and obtain

$$\begin{aligned} |\{f \in X_r(N) \mid f \text{ is reducible}\}| &\leq (\sqrt{2N+1} + L)^{2r} H^{-1} \\ &= (2N+1)^r \left(1 + \frac{1}{r}\right)^{2r} H^{-1} \\ &\ll r^2(2N+1)^{r-1/2}(\log(2N+1)), \end{aligned}$$

² One knows explicit lower bounds $\pi(L) \geq \alpha L(\log L)^{-1}$ for $L \geq 2$.

and, on the other hand, if $\sqrt{2N+1} < \alpha r^2(\log 2r)$, then this estimate is trivial, provided the implied constant is chosen to be big enough. So we have

$$|\{f \in X_r(N) \mid f \text{ is reducible}\}| \ll r^2(2N+1)^{r-1/2}(\log N)$$

for $N \geq 2$ and $r \geq 1$, with an absolute implied constant.

If we are interested rather in the splitting field, we must refine the sieve argument. Assume $f \in \mathbf{Z}[T]$ is an irreducible monic polynomial of degree r , and let K_f be its splitting field, generated over \mathbf{Q} by the r distinct roots $\Theta = (\theta_1, \dots, \theta_r)$ of f . The Galois group $G = \text{Gal}(K_f/\mathbf{Q})$ is the group of field automorphisms of K_f , and its action on the finite set Θ gives an embedding of G in the permutation group of Θ , which is of course isomorphic to the symmetric group \mathfrak{S}_r on r letters. This isomorphism is not canonical, but in what follows this is not a problem.

Now what is needed is the fact that the factorization of the reduction of f modulo a prime ℓ gives information on the group G , or rather on the image \tilde{G} of G inside \mathfrak{S}_r . Indeed, if f factors in $\mathbf{F}_\ell[T]$ without square factors (as it will for all but finitely many ℓ), then we can write this factorization in the form

$$f \pmod{\ell} = f_1 \cdots f_r \in \mathbf{F}_\ell[T], \quad (4.2)$$

where each f_i is a products of $n_i \geq 0$ distinct irreducible monic polynomials in $\mathbf{F}_\ell[T]$ of degree i , so that

$$n_1 + 2n_2 + \cdots + rn_r = r,$$

and then, it is well known (going back to Dedekind at least) that the image \tilde{G} of G in the symmetric group contains an element σ with cycle-type described by n_1 fixed points, n_2 disjoint transpositions, \dots , n_k disjoint cycles of length k , \dots . Indeed, this element is obtained by ‘lifting’ the Frobenius automorphism $x \mapsto x^\ell$ acting on a finite splitting field of f modulo ℓ : this automorphism acts cyclically on the i roots of each factor of degree i of f_i of f , so it has the right cycle structure, and there remains to see that it can be obtained by reduction from an element of the Galois group of the splitting field of f over \mathbf{Q} ; see, e.g., [86, IX, Theorem 2.9] for the latter.

This is sufficient to determine the Galois group of the splitting field, because of a classical group-theoretic result: in a finite group G , there is no proper subgroup H which contains an element of every conjugacy class in G .³ For any conjugacy class c of \mathfrak{S}_r described by permutations with a given cycle type

³ Indeed, if H is such a subgroup, since there are at most $|G/H|$ distinct conjugates of H , their union has to be disjoint in order to cover G , which means that there is a single conjugate since they are subgroups, and then $H = G$.

(described as elements which are products of n_i disjoint i -cycles for $1 \leq i \leq r$), define $\Omega_{c,\ell}$ to be the set of monic polynomials of degree r in $\mathbf{F}_\ell[T]$ which factor as in (4.2); we can therefore write

$$E_r(N) \subset \bigcup_{c \in \mathfrak{S}_r^\#} S(X_r(N), \Omega_c; \mathcal{L}^*)$$

where \mathcal{L}^* is again an arbitrary finite set of primes. With \mathcal{L}^* the set of primes $\ell \leq L$ as before, and $\mathcal{L} = \mathcal{L}^*$, and if r is considered fixed, we need only apply (asymptotically) Lemma B.1 to obtain the lower bound for H that implies

$$S(X_r(N), \Omega_c; \mathcal{L}^*) \ll N^{r-1/2}(\log N)$$

for any c and

$$E_r(N) \ll N^{r-1/2}(\log N).$$

To deal with uniformity in r , the main issue is in fact the size of the density factor

$$M = \sum_c \delta_c^{-1}, \quad \text{where} \quad \frac{|\Omega_{c,\ell}|}{\ell^r} \sim \delta_c, \text{ as } \ell \rightarrow +\infty;$$

indeed, the argument above shows that in a uniform estimate with respect to r , the right-hand side will involve

$$M(\sqrt{2N+1} + L)^{2r} L^{-1}(\log L).$$

As stated, this gives a very poor dependency on r because $\delta_1 = 1/r!$ is very small and $M \geq \delta_1^{-1}$. However, it is possible to judiciously select sets $C \subset \mathfrak{S}_r$ of conjugacy classes with much larger density so that no proper subgroup of \mathfrak{S}_r contains an element of each class $c \in C$.

Indeed, Gallagher [46, Lemma, p. 98] quotes a lemma of Bauer (with a very cute proof by D. Knutson) to the effect that no proper subgroup of \mathfrak{S}_r acting transitively on $\{1, \dots, r\}$ contains both a transposition and a cycle of prime order $p > r/2$. Since the Galois group of the splitting field of a polynomial f acts transitively on the roots if and only if the polynomial is irreducible, this means we can take $C = C_1 \cup C_2$ where C_1 is union of conjugacy classes of elements with a single transposition and products of cycles of odd lengths, and C_2 is the union of conjugacy classes of elements of prime order $p > r/2$, and we then see that

$$E_r(N) \subset \{f \in X_r(N) \mid f \text{ is reducible}\} \cup S(X_r(N), \Omega^1; \mathcal{L}^*) \cup S(X_r(N), \Omega^2; \mathcal{L}^*)$$

with Ω_ℓ^i the set of monic polynomials of degree r in $\mathbf{F}_\ell[T]$ which factor as prescribed by the cycle type of an element of C_i .

Gallagher [46, p. 99] shows that elements in C_1 and C_2 have density

$$\frac{|\{\sigma \in C_1\}|}{r!} \sim \frac{\log 2}{\log r}, \quad \text{as } r \rightarrow +\infty, \quad (4.3)$$

$$\frac{|\{\sigma \in C_2\}|}{r!} \sim \frac{1}{\sqrt{2\pi r}}, \quad \text{as } r \rightarrow +\infty. \quad (4.4)$$

Using this, Lemma B.1 and the mean-value theorem in a way similar to the estimation of the number of reducible elements of $X_r(N)$ above, one is led to

$$E_r(N) \ll r^3(2N+1)^{r-1/2}(\log N)$$

for $N \geq 2$ and $r \geq 1$, with an absolute implied constant. \square

Remark 4.3 To indicate the flexibility and versatility of the general theory, here is a different way to set up a sieve to tackle this problem: take $X = Y$ to be the set of monic polynomials of degree r in $\mathbf{Z}[X]$ with coefficients $\leq N$ in absolute value (with F being the identity and μ the counting measure); then for all primes ℓ , let Y_ℓ be the set of conjugacy classes in \mathfrak{S}_r , and define $\rho_\ell(P)$ to be the conjugacy class associated with the cycle type giving the factorization of P modulo ℓ (ignoring, for simplicity, the issue of primes dividing the discriminant), so that the Galois group of the splitting field of P , as subgroup of \mathfrak{S}_r , contains an element in the conjugacy class $\rho_\ell(P)$ for all ℓ . Then the set of polynomials with small splitting field satisfies

$$E_r(N) = \bigcup_{\sigma^\sharp \in \mathfrak{S}_r^\sharp} S(X, \{\sigma^\sharp\}; \{\text{primes}\})$$

and can (or could) therefore be estimated using this sieve setting. However, it is very unclear (perhaps unlikely) that the large sieve constant can be estimated well enough in this situation to recover (or improve on) Gallagher's Theorem.

Remark 4.4 Another important remark arises from this proof, which illustrates a fairly general point concerning the large sieve: Gallagher's Theorem (for fixed r) should not really be thought of as an *existence* theorem for irreducible polynomials, or polynomials with large Galois group, and not even a 'full density' result (in the sense that $|E_r(N)|/|X_r(N)| \rightarrow 0$ as $N \rightarrow +\infty$). If such a result, and no more, was the goal, then it would have sufficed⁴ to know the density of polynomials over a large finite field with given splitting type to obtain a bound $\leq (1 - \delta)$, for some $\delta > 0$, for $\limsup |E_r(N)|/|X_r(N)|$, and then to repeat with k distinct primes to replace this by $(1 - \delta)^k$ for arbitrary k . And of

⁴ This is the idea of van der Waerden.

course, all ingredients for such a proof are essentially necessary ingredients for applying the large sieve.

On the other hand, there *are* applications where a quantitative bound is desired, and moreover the control of the uniformity of the estimates over r is also interesting; then the large sieve is perfectly suited for the task, where the other techniques fail (or become too unwieldy to be pursued; there is nothing inherently ineffective in van der Waerden's argument). Also, the experience from classical problems of analytic number theory is that, when the large sieve is merely one tool among others to solve a problem, then it is really required in its full power, and can not be replaced with essentially weaker arguments.

4.3 The multiplicative large sieve inequality

In the historical development of the large sieve, an important role was played by the derivation from (4.1) of a similar inequality involving *multiplicative* Dirichlet characters, which was a key ingredient in the proof of the Bombieri–Vinogradov Theorem (see, e.g., [67, Chapter 17] for the latter).

To state this inequality, recall first that a Dirichlet character χ modulo $q \geq 1$ is an arithmetic function

$$\chi : \mathbf{Z} \rightarrow \mathbf{C}$$

defined as the composite

$$\mathbf{Z} \rightarrow \mathbf{Z}/q\mathbf{Z} \xrightarrow{\chi_0} \mathbf{C}$$

where χ_0 is the extension by zero to $\mathbf{Z}/q\mathbf{Z}$ of a homomorphism of $(\mathbf{Z}/q\mathbf{Z})^\times$ to \mathbf{C}^\times . An example is the *trivial* character ε_q obtained when χ_0 is the trivial character. Note that ε_q is not equal identically to 1, except for $q = 1$, since $\varepsilon_q(n) = 0$ if $(n, q) \neq 1$.

Because of the Chinese Remainder Theorem, a Dirichlet character modulo q can be expressed uniquely as follows

$$\chi(n) = \prod_{p|q} \chi_p(n)$$

where χ_p is a Dirichlet character modulo $p^{v_p(q)}$, $v_p(q)$ being the power of p dividing q . If none of the characters χ_p is the trivial character (modulo $p^{v_p(q)}$), the character χ is said to be *primitive*, and q is its *conductor*.

Now the multiplicative large sieve inequality is the following:

$$\sum_{q \leq Q} \sum_{\chi \pmod{q}}^* \left| \sum_{M < n \leq M+N} a_n \chi(n) \right|^2 \leq \Delta \sum_{M < n \leq M+N} |a_n|^2, \quad (4.5)$$

for arbitrary complex numbers (a_n) , where the sum over χ is restricted to primitive characters with conductor q .

Exercise 4.1 Deduce (4.5) with $\Delta \leq N - 1 + Q^2$ from the inequality in Theorem 4.1 by proving the identity

$$q \sum_{\chi \pmod{q}}^* \left| \sum_{M < n \leq M+N} a_n \chi(n) \right|^2 = \varphi(q) \sum_{\substack{a \pmod{q} \\ (a, q) = 1}} \left| \sum_{M < n \leq M+N} a_n e\left(\frac{an}{q}\right) \right|^2$$

for any q (see, e.g., [67, Theorem 7.13] for details).

Can this inequality be related to our general setting? Indeed, in the following way. Let $L \geq 2$ be given, and let Y be the set of integers $n \in \mathbf{Z}$ not divisible by primes $\leq L$. Then taking Λ to be the set of primes $\ell \leq L$, the restriction to Y of the reduction map modulo ℓ gives a surjection

$$Y \rightarrow (\mathbf{Z}/\ell\mathbf{Z})^\times = Y_\ell$$

for any such ℓ , and moreover it is natural to use the density

$$v_\ell(y) = \frac{1}{\ell - 1}$$

for $\ell \leq L$. If we take X to be the elements in Y which are $\leq N$, with $F_x = x$, and \mathcal{L}^* the set of primes $\ell \leq L$, the sifted sets become

$$S(X, \Omega; \mathcal{L}^*) = \{n \leq N \mid (p \mid n \Rightarrow p > L) \text{ and } n \pmod{\ell} \notin \Omega_\ell \text{ for } \ell \leq L\},$$

where $\Omega_\ell \subset (\mathbf{Z}/\ell\mathbf{Z})^\times$. Note that, in particular, $S(X, \Omega; \mathcal{L}^*)$ contains the set of primes p with $L < p \leq N$ and $p \pmod{\ell} \notin \Omega_\ell$ for $\ell \leq L$.

Take for \mathcal{L} the set of squarefree numbers $\leq L$. A simple check (using the definition of primitive characters above) shows that Y_m , for $m \in \mathcal{L}$, is naturally identified with $(\mathbf{Z}/m\mathbf{Z})^\times$, and the set of primitive characters modulo m can be chosen as the basis \mathcal{B}_m^* of $L^2(Y_m)$. Hence the inequality defining the large sieve constant Δ becomes

$$\sum_{m \leq L}^b \sum_{\chi \pmod{m}}^* \left| \sum_{n \in X} a_n \chi(n) \right|^2 \leq \Delta \sum_{n \in X} |a_n|^2$$

for any complex numbers a_n , where χ runs over primitive characters modulo m ; note that when $n \in X$, we have $\chi(n) \neq 0$ since $(n, m) = 1$ by definition. It follows that $\Delta \leq N - 1 + L^2$ by the multiplicative large sieve inequality (4.5).

In particular, we have

$$|S(X, \Omega; \mathcal{L}^*)| \leq \Delta H^{-1}$$

where

$$H = \sum_{m \leq L} \prod_{\ell | m} \frac{|\Omega_\ell|}{\ell - 1 - |\Omega_\ell|}$$

and $\Delta \leq N - 1 + L^2$ is the multiplicative large sieve constant.

4.4 The elliptic sieve

The next application is an (apparently) new use of the classical large sieve. Let E/\mathbf{Q} be an elliptic curve defined over \mathbf{Q} , given by an affine Weierstrass equation

$$y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6, \quad \text{where } a_i \in \mathbf{Z}. \quad (4.6)$$

The set $E(\mathbf{C})$ of complex points of E (in other words, the set of solutions $(x, y) \in \mathbf{C} \times \mathbf{C}$ of this equation, with the addition of the point at infinity obtained when taking the closure in the projective plane) has a well-known structure of abelian group, with the point at infinity being the origin, the group law being further dictated, essentially, by the condition that three distinct points sum to zero if and only if they are collinear. For this and other basic facts about elliptic curves, we refer to Silverman's classic book [124].

We will consider as object to sieve the Mordell–Weil group $E(\mathbf{Q})$, i.e., the subgroup of $E(\mathbf{C})$ consisting of those points which have rational coordinates, together with the point at infinity. Because the parameters a_i are integers, this is indeed a subgroup of $E(\mathbf{C})$, and the famous Mordell–Weil theorem (due to Mordell in this special case) states that this is an abelian group with finite rank.

Now let Λ_E be the set of primes ℓ of good reduction (or the possibly smaller set of primes not dividing the discriminant of E). For $\ell \in \Lambda_E$, we can define the reduction of E modulo ℓ as the elliptic curve over \mathbf{F}_ℓ given by the same equation (4.6), where a_i are reduced modulo ℓ . In particular, the set of solutions in \mathbf{F}_ℓ , with the point at infinity, which we denote $E(\mathbf{F}_\ell)$, is a finite abelian group. Crucially, we have a well-defined reduction homomorphism

$$\rho_\ell : E(\mathbf{Q}) \rightarrow E(\mathbf{F}_\ell)$$

obtained in a fairly obvious manner: if ℓ divides either of the denominators of the coordinates of a point $x \in E(\mathbf{Q})$, we map x to the point at infinity, and otherwise each coordinate is mapped to an element of \mathbf{F}_ℓ by inverting the denominator in \mathbf{F}_ℓ (see, e.g., [124, VII.2] for more details).

In general, this map ρ_ℓ is not onto (indeed, quite often, $E(\mathbf{Q})$ is finite, whereas the Hasse inequality states that $|E(\mathbf{F}_\ell)| = \ell + 1 - a_E(\ell)$ with $|a_E(\ell)| \leq 2\sqrt{\ell}$; see, e.g., [124, V.1]). In order to have a sieve setting in our sense, we should define E_ℓ to be the *image* of $E(\mathbf{Q})$ in $E(\mathbf{F}_\ell)$, hoping that no confusion will arise

from the similarity in notation. Thus we are given an interesting-looking sieve setting $(E(\mathbf{Q}), \Lambda_E, (\rho_\ell : E(\mathbf{Q}) \rightarrow E_\ell))$.

Let us now consider what siftable sets are suggested by the arithmetic of elliptic curves. First of all, if $E(\mathbf{Q})$ is a finite group (the rank of E is zero), there does not seem much to be said, especially since Mazur has shown that only finitely many groups can arise as $E(\mathbf{Q})$, and that there exist good algorithms to find which holds for a given curve (see, e.g., [124, VIII.7] for more information).

So assume that the rank of $E(\mathbf{Q})$ is at least one. Then the most natural sets $X \subset E(\mathbf{Q})$ for sieving are the finite sets of rational points $x \in E(\mathbf{Q})$ with ‘height’ bounded by some T . The height may be calculated by means of the naïve height h_n , defined by $h_n(0) = 0$ for the point at infinity, and

$$h_n(x) = \log \max(|p|, |q|)$$

for $x = (p/q, r/s)$, where (p, q, r, s) are integers, $(p, q) = (r, s) = 1$. For many purposes, it is better to consider the canonical height h , which is a map

$$h : E(\mathbf{Q}) \rightarrow [0, +\infty[$$

with the following properties: (1) $h(x) = 0$ if and only if x is of finite order; (2) on $E(\mathbf{Q}) \otimes \mathbf{R}$, h is a positive definite quadratic form; (3) we have

$$h(x) = \frac{1}{2}h_n(x) + O(1)$$

for $x \in E(\mathbf{Q})$. In our results, both heights give therefore equivalent results; see, e.g., [124, VIII.4, VIII.9], in particular Theorem VIII.9.3 of [124].

There is some interest in sieving $E(\mathbf{Q})$ because of the following property of the reduction map which has already been described: a rational point $x = (r, s) \in E(\mathbf{Q})$ (in affine coordinates, so x is non-zero in $E(\mathbf{Q})$) maps to a non-zero point $E(\mathbf{F}_\ell)$ if and only if ℓ does not divide the *denominator* of the affine coordinates r and s of the point. In particular, *integral solutions* of (4.6) are elements which remain after sieving by $\Omega_\ell = \{0\}$ for *all* $\ell \in \Lambda_E$. It is well known (Siegel’s Theorem) that there are only finitely many such integral solutions. However, we will not try to use sieve to investigate integral points; rather, we will use below a strengthening of Siegel’s Theorem to show that a suitable subset of Λ_E gives a sieve which closely resembles the classical sieve of integers modulo primes.

We now use these ideas to prove Theorem 1.1, showing that most rational points have denominators divisible by many (small) primes. First, define $\omega_E(x)$ to be the number of prime factors dividing the denominator of the coordinates of x , without multiplicity, with $\omega_E(0) = +\infty$. We recall the statement:

Theorem 4.5 *Let E/\mathbf{Q} be an elliptic curve as above. Assume that the rank r of $E(\mathbf{Q})$ is $r \geq 1$. Then we have*

$$|\{x \in E(\mathbf{Q}) \mid h(x) \leq T\}| \sim c_E T^{r/2} \quad (4.7)$$

as $T \rightarrow +\infty$, for some constant $c_E > 0$; and for any fixed real number κ with $0 < \kappa < 1$, we have

$$|\{x \in E(\mathbf{Q}) \mid h(x) \leq T \text{ and } \omega_E(x) < \kappa \log \log T\}| \ll T^{r/2} (\log \log T)^{-1},$$

for $T \geq 3$, where the implied constant depends only on E and κ .

Proof Since $E(\mathbf{Q})$ is of finite rank, we can find a free subgroup $M \simeq \mathbf{Z}^r$ of $E(\mathbf{Q})$ such that

$$E(\mathbf{Q}) = M \oplus E(\mathbf{Q})_{\text{tors}}, \quad \text{with } E(\mathbf{Q})_{\text{tors}} \text{ finite.}$$

Then let (x_1, \dots, x_r) be a fixed \mathbf{Z} -basis of M , and let M' be the subgroup generated by (x_2, \dots, x_r) . We will perform sieving only on affine ‘lines’ directed by x_1 , passing through a point of M' .

But first of all, since the canonical height is a positive definite quadratic form on $E(\mathbf{Q}) \otimes \mathbf{R} = M \otimes \mathbf{R}$, the asymptotic formula (4.7) is clear:⁵ it amounts to nothing else but counting integral points in $M \otimes \mathbf{R} \simeq \mathbf{R}^r$ with norm $\sqrt{h(x)} \leq \sqrt{T}$, this being repeated as many times as there are torsion cosets.

For convenience, we will now measure the size of elements in $E(\mathbf{Q})$ using the squared L^∞ -norm:

$$\|x\|_\infty^2 = \max |a_i|^2, \quad \text{for } x = \sum a_i x_i + t \text{ with } t \in E(\mathbf{Q})_{\text{tors}};$$

this satisfies $h(x) \asymp \|x\|_\infty^2$ for all $x \in M$, the implied constants depending only on E , simply by comparison of two norms on a finite-dimensional \mathbf{R} -vector space.

Now the actual sieve result is the following:

Lemma 4.6 *For any fixed $\kappa \in]0, 1[$, any fixed $x' \in M'$, any fixed torsion point $t \in E(\mathbf{Q})_{\text{tors}}$, we have*

$$|\{x \in (t+x') \oplus \mathbf{Z}x_1 \mid \|x\|_\infty^2 \leq T \text{ and } \omega_E(x) < \kappa \log \log T\}| \ll \sqrt{T} (\log \log T)^{-1},$$

for $T \geq 3$, the implied constant depending only on E , κ and x_1 , but not on x' or t .

⁵ And of course it is not new, but is included in the statement in order to clarify the gain in the estimate for the number of points with denominators involving few primes.

Taking this for granted, we conclude immediately that

$$|\{x \in E(\mathbf{Q}) \mid h(x) \leq T \text{ and } \omega_E(x) < \kappa \log \log T\}| \ll T^{r/2}(\log \log T)^{-1},$$

by summing the inequality of the lemma over all $x' \in M'$ with $\|x'\|_\infty^2 \leq T$ and over all $t \in E(\mathbf{Q})_{tors}$ (the number of pairs (t, x') is $\ll T^{(r-1)/2}$), the implied constant depending only on E and the choice of basis of M . \square

To prove Lemma 4.6, the crucial tool is the following result which makes the link between our sieve and the diophantine properties of S -integral points on elliptic curves.

Lemma 4.7 *Let x_1 be a point of infinite order in $E(\mathbf{Q})$. For $\ell \in \Lambda_E$, let $\nu(\ell)$ be the order of x_1 modulo ℓ in the finite group $E(\mathbf{F}_\ell)$. Then all but finitely many primes p occur as the value of $\nu(\ell)$ for some ℓ of good reduction.*

Proof For a prime p , consider $px_1 \in E(\mathbf{Q})$. A prime ℓ of good reduction divides the denominator of the coordinates of px_1 if and only if $p \equiv 0 \pmod{\nu(\ell)}$, which means that $\nu(\ell)$ is either 1 or p . So if p is not of the form $\nu(\ell)$, it follows that px_1 is an S -integral point of $E(\mathbf{Q})$,⁶ where S is the union of the set of primes of bad reduction and the finite set of primes where $x_1 \equiv 0 \pmod{\ell}$ (the latter is finite because only finitely many primes divide the denominator of the coordinates of $x_1 \neq 0$). By Siegel's finiteness theorem (see, e.g., [124, Theorem IX.4.3]), there are only finitely many S -integral solutions to (4.6), hence finitely many possibilities for px_1 for such p ; because x_1 is assumed to be of infinite order, this translates to finitely many p which are not of the form $\nu(\ell)$. \square

Note that this lemma is also a trivial consequence of a result of Silverman [122, Proposition 10] according to which all but finitely many *integers* are of the form $\nu(\ell)$ for some ℓ . The proofs are indeed related, since Silverman's result depends on a stronger form of Siegel's theorem.

Proof of Lemma 4.6 Fix $x' \in M'$, $t \in E(\mathbf{Q})_{tors}$. The left-hand side of the lemma being zero unless $\|t + x'\|_\infty^2 \leq T$, we assume that this is the case. We will use the following group sieve setting:

$$\begin{aligned} \Psi &= (\mathbf{Z}x_1, \Lambda_E, \mathbf{Z}x_1 \rightarrow \rho_\ell(\mathbf{Z}x_1) \subset \rho_\ell(E(\mathbf{Q}))), \\ X &= \{mx_1 \in G \mid \|t + x' + mx_1\|_\infty^2 = m^2 \leq T\}, \quad F_x = x. \end{aligned}$$

⁶ Recall that an S -integer is a rational number with (minimal) denominator only divisible by primes in S , and an S -integral solution of (4.6) is one where both coordinates are S -integers.

For any prime $\ell \in \Lambda_E$, the finite group G_ℓ is a quotient of $\mathbf{Z}x_1$ and is isomorphic to $\mathbf{Z}/v(\ell)\mathbf{Z}$ where $v(\ell)$ is the order of the reduction of x_1 modulo ℓ . So this sieve is really an ordinary-looking one for integers, except for the use of reductions modulo $v(\ell)$ instead of reductions modulo primes.

We select the prime sieve support $\mathcal{L}^* \subset \Lambda_E$ containing all ℓ for which $v(\ell)$ is a prime number $p \leq L$, where, if the same prime p occurs as values of $v(\ell)$ for two or more primes, we keep only one, and we also put $\mathcal{L} = \mathcal{L}^*$.

The point is that the inequality defining the large sieve constant here is

$$\sum_{\ell \in \mathcal{L}} \sum_{a \pmod{v(\ell)}}^* \left| \sum_{|m| \leq \sqrt{T}} \alpha(m) e\left(\frac{am}{v(\ell)}\right) \right|^2 \leq \Delta \sum_{|m| \leq \sqrt{T}} |\alpha(m)|^2, \quad (4.8)$$

for all $(\alpha(m))$, and this may be reformulated as

$$\sum_{p \leq L}^* \sum_{a \pmod{p}}^* \left| \sum_{|m| \leq \sqrt{T}} \alpha(m) e\left(\frac{am}{p}\right) \right|^2 \leq \Delta \sum_{|m| \leq \sqrt{T}} |\alpha(m)|^2.$$

where \sum^* in the sum over p indicates that only those p which occur as $v(\ell)$ for some ℓ are taken into account. We recognize a subsum of the classical large sieve inequality, and by positivity, it follows that

$$\Delta \leq 2\sqrt{T} + L^2$$

for $L \geq 2$. We now apply Proposition 2.15: we have

$$\sum_{x \in X} \left(P(x, \mathcal{L}) - P(\mathcal{L}) \right)^2 \leq \Delta Q(\mathcal{L}) \quad (4.9)$$

where $P(x, \mathcal{L})$, $P(\mathcal{L})$ and $Q(\mathcal{L})$ are defined in (2.14), (2.15) for any given choice of sets $\Omega_\ell \subset G_\ell$ for $\ell \in \Lambda_E$.

We let $\Omega_\ell = \{-\rho_\ell(t + x')\}$. By the remark before the statement of Theorem 4.5, we have $\rho_\ell(mx_1) \in \Omega_\ell$ if and only if ℓ divides the denominator of the coordinates of $t + x' + mx_1$, and therefore for $x = mx_1 \in X$, we have

$$P(mx_1, \mathcal{L}) \leq \omega_E(t + x' + mx_1).$$

On the other hand, we have

$$P(\mathcal{L}) = \sum_{\ell \in \mathcal{L}} \frac{1}{|G_\ell|} = \sum_{\ell \in \mathcal{L}} \frac{1}{v(\ell)} = \sum_{p \leq L} \frac{1}{p} + O(1) = \log \log L + O(1)$$

for any $L \geq 3$, because, by Lemma 4.7, the values $v(\ell) \leq L$ range over all primes $\leq L$, with only finitely many exceptions.

Hence there exists L_0 depending on E , x_1 and κ only, such that if $L \geq L_0$, we have

$$P(\mathcal{L}) \geq \frac{1+\kappa}{2} \log \log L.$$

Putting together these two inequalities, we see that if we assume $T \leq L^2$, say, and $L \geq L'_0$ for some other constant L'_0 (depending on E , x_1 and κ), then for any $mx_1 \in X$ such that $t+x'+mx_1$ satisfies $\omega_E(t+x'+mx_1) < \kappa \log \log T$, we have

$$\left(P(x, \mathcal{L}) - P(\mathcal{L})\right)^2 \gg (\log \log T)^2,$$

the implied constant depending only on E , x_1 and κ . So it follows by positivity from (4.9) and the inequality $Q(\mathcal{L}) \leq P(\mathcal{L}) \ll \log \log T$ that

$$\begin{aligned} |\{x \in t + x' \oplus \mathbf{Z}x_1 \mid \|x\|_\infty^2 \leq T \text{ and } \omega_E(x) < \kappa \log \log T\}| \\ \ll \frac{\Delta}{\log \log T} \ll \frac{\sqrt{T} + L^2}{\log \log T} \end{aligned}$$

for any $L \geq L'_0$. If $T^{1/2} \geq L'_0$, we take $L = T^{1/2}$ and prove the inequality of the lemma directly, and otherwise we need only increase the resulting implied constant to make it valid for all $T \geq 3$, since L'_0 depends only on E , x_1 and κ . \square

Exercise 4.2 Prove the following analogue of Theorem 4.5 and Lemma 4.6 for the multiplicative group instead of an elliptic curve: show that for any integer $a \notin \{\pm 1\}$, and for any fixed real number $\kappa \in]0, 1[$, we have

$$|\{n \leq N \mid \omega(a^n - 1) < \kappa \log \log N\}| \ll \frac{N}{\log \log N}$$

for all $N \geq 3$, the implied constant depending only on a and κ . [Hint: The analogue of Lemma 4.7 may be obtained either by invoking the finiteness of the number of solutions to S -unit equations, or a theorem of Schinzel which is the exact analogue of Silverman's theorem in [122].]

Notice the similarity between the above discussion and the Hardy–Ramanujan results concerning the normal order of the number of prime divisors of an integer (see, e.g., [56, 22.11]), in Turán's formulation (see (2.16)).

However, our result does not imply that the normal order of $\omega_E(x)$ for x with $h(x) \leq T$ is $\log \log T$, because the definition of the height is logarithmic in terms of the denominator of the coordinates of x , so that we can expect that the

denominators of rational points x are typically of size $\exp h(x)$. As such, they should have about

$$\log \log \exp(h(x)) = \log(h(x)) \sim \log T$$

prime divisors in order to be ‘typical’ integers. Yet, one may notice that the proof of the theorem really produces many *small* prime factors: indeed, in our sieve we detect only prime factors ℓ where x_1 has order $\leq T^{1/2}$ modulo ℓ , which we may expect to be mostly primes of size T (in logarithmic scale), hence of size $\log h(x)$. Now it is typical behaviour for an integer n of size X (here $h(x) \leq T$) to have roughly $\log \log \log X$ prime divisors of this size (here about $\log \log T$), although this is at the border of more unpredictable behaviour; such results are due in particular to Erdős (see p. 135 in Ruzsa’s survey [110]). It would be very interesting to know whether the distribution of $\omega_E(x)$ is as regular as $\omega(n)$. Indeed, it would be interesting to know this for $\omega(a^n - 1)$, as in Exercise 4.2.

Note also that, as mentioned during the discussion of Proposition 2.15, applying the (apparently stronger) form of the large sieve involving square-free numbers would only give a bound for the number of points which are \mathcal{L} -integral. Since (for any finite set S), there are only finitely many S -integral points, and moreover this is used in the proof of Lemma 4.7, this would not be a very interesting conclusion.

We conclude by relating this sieve, more precisely Lemma 4.6, to so-called *elliptic divisibility sequences*, a notion introduced by M. Ward and currently the subject of a number of investigations by Ayad, Silverman, T. Ward, Everest, and others (see, e.g., [6], [123], [130], [36]). This shows that the proposition above has very concrete interpretations.

Proposition 4.8 *Let $(W_n)_{n \geq 0}$ be an unbounded sequence of integers such that*

$$\begin{aligned} W_0 &= 0, \quad W_1 = 1, \quad W_2 W_3 \neq 0, \quad W_2 \mid W_4, \\ W_{m+n} W_{m-n} &= W_{m+1} W_{m-1} W_n^2 - W_{n+1} W_{n-1} W_m^2, \quad \text{for } m \geq n \geq 1, \\ \Delta &= W_4 W_2^{15} - W_3^3 W_2^{12} + 3 W_4^2 W_2^{10} - 20 W_4 W_3^3 W_2^7 \\ &\quad + 4 W_4^3 W_2^5 + 16 W_3^6 W_2^4 + 8 W_4^2 W_3^3 W_2^2 + W_4^4 \neq 0. \end{aligned}$$

Then for any κ such that $0 < \kappa < 1$, we have

$$|\{n \leq N \mid \omega(W_n) < \kappa \log \log N\}| \ll \frac{N}{\log \log N}$$

for $N \geq 3$, where the implied constant depends only on κ and (W_n) .

Proof This depends on the relation between elliptic divisibility sequences and pairs (E, x_1) of an elliptic curve E/\mathbf{Q} and a point $x_1 \in E(\mathbf{Q})$. Precisely (see, e.g. [36, Section 2]) there exists such a pair (E, x_1) with x_1 of infinite order such that if $(a_n), (b_n), (d_n)$ are the (unique) sequences of integers with $d_n \geq 1$, $(a_n, d_n) = (b_n, d_n) = 1$ and

$$nx_1 = \left(\frac{a_n}{d_n^2}, \frac{b_n}{d_n^3} \right),$$

then we have

$$d_n \mid W_n \text{ for } n \geq 1$$

(without the condition $\Delta = 0$, this is still true provided *singular* elliptic curves are permitted; the condition that (W_n) be unbounded implies that x_1 is of infinite order).

Now the primes dividing d_n are precisely those dividing the denominators of the coordinates of the points in $\mathbf{Z}x_1$, and we have therefore

$$\omega(W_n) \geq \omega(d_n) = \omega_E(nx_1).$$

Hence Lemma 4.6 gives the desired result. \square

The ‘simplest’ example is the sequence (W_n) given by

$$W_0 = 0, \quad W_1 = 1, \quad W_2 = 1, \quad W_3 = -1, \quad W_4 = 1,$$

$$W_n = \frac{W_{n-1}W_{n-3} + W_{n-2}^2}{W_{n-4}}, \quad \text{for } n \geq 4$$

(sequence A006769 in the Online Encyclopedia of Integer Sequences, www.research.att.com/~njas/sequences/), which corresponds to the case of $E : y^2 - y = x^3 - x$ and $x_1 = (0, 0)$.

Elliptic divisibility sequences are natural generalizations of non-degenerate divisibility sequences (u_n) defined by linear recurrence relations of order 2, the simplest of which are $u_n = a^n - 1$ where $a \geq 2$ is an integer. The result of Exercise 4.2 clearly shows the analogy. It would actually be more interesting to find a result showing a *difference* between this case and the case of elliptic divisibility sequences (this is expected, e.g., because

$$\sum_n \frac{1}{\log(a^n - 1)} = +\infty, \quad \sum_n \frac{1}{\log W_n} < +\infty,$$

the latter because $\log W_n$ is about the same as $\log d_n \asymp \log \exp(h(nx)) \asymp n^2$ since h is a quadratic form, and those series are heuristically the expected number of primes in the sequences considered).

4.5 Other examples

We now list, without details, some interesting variants of the large sieve.

Example 4.9 Serre [117] has used a variant of the higher-dimensional large sieve where

$$\Psi = (\mathbf{Z}^r, \{\text{primes}\}, \mathbf{Z}^r \rightarrow (\mathbf{Z}/\ell^2\mathbf{Z})^r)$$

and

$$X = \{(x_1, \dots, x_r) \in \mathbf{Z}^r \mid |x_i| \leq N\}$$

with $F_x = x$. With suitable sieving sets, this provides estimates for the number of trivial specializations of elements of 2-torsion in the Brauer group of $\mathbf{Q}(T_1, \dots, T_r)$.

Example 4.10 Here is a new example, which is a number field analogue of the situation of [80] (described also in Chapter 8). It is related to Serre's discussion in [116] of a higher-dimensional Chebotarev density theorem over number fields (see also [104] for an independent treatment with more details). Let Y/\mathbf{Z} be a separated scheme of finite type, and let $Y_\ell \rightarrow Y$ be a family of étale Galois coverings,⁷ corresponding to surjective maps $G = \pi_1(Y, \bar{\eta}) \rightarrow G_\ell$. The sieve setting is $(G, \{\text{primes}\}, G \rightarrow G_\ell)$. Now let $|Y|$ denote the set of closed points of Y , which means those where the residue field $k(y)$ is finite, and let

$$X = \{y \in |Y| \mid |k(y)| \leq T\}$$

for some $T \geq 2$, which is finite. For $y \in X$, denote by $F_x \in G$ the corresponding geometric Frobenius automorphism (or conjugacy class rather) to obtain a siftable set $(X, \text{counting measure}, F)$ associated with the conjugacy sieve. It should be possible to obtain a large sieve inequality in this context, at least assuming the Generalized Riemann Hypothesis and the Artin conjecture.

Note that if Y is the set of prime ideals in the ring of integers in some number field (i.e., the spectrum of some such ring, even for $Y = \mathbf{Z}$ itself), this becomes the sieve for Frobenius considered by D. Zywna ('The large sieve and Galois representations', preprint), with conditional applications to the Lang–Trotter conjecture, and to Koblitz's conjecture for elliptic curves over number fields.

Example 4.11 In [105], Poonen uses a 'closed-point sieve' to study (among other things) the homogeneous polynomials $f \in \mathbf{F}_q[x_0, \dots, x_n]$ for which the intersection $Z \cap \{f = 0\}$ is smooth, where $Z \subset \mathbf{P}^n$ is a fixed smooth

⁷ Or better with 'controlled' ramification, if not étale, since this is likely to be needed for some natural applications.

(quasi)projective variety defined over \mathbf{F}_q . This can be phrased roughly as follows in our terminology: Y is the set of non-zero homogeneous polynomials in $\mathbf{F}_q[x_0, \dots, x_n]$, Λ is the set of closed points of Z , and for any such point $x \in \Lambda$, we define ρ_x and Y_x by considering the image of the natural \mathbf{F}_q -linear map $\rho_x : Y \rightarrow V_x$ where V_x is the vector space of Taylor expansions of order 1 at x (over the residue field of x , but seen as \mathbf{F}_q -vector space). Note that Y_x is usually far from being the whole vector space V_x because elements of Y have coefficients in \mathbf{F}_q , not in the residue field of x (see [105, Section 2.1] for the more intrinsic characterization of Y_x as \mathbf{F}_q -vector space of global sections on some finite subscheme of Z , when d is large enough).

This defines the sieve setting $(Y, \Lambda, (\rho_x))$, and the relevant siftable sets are the finite sets $X = X_d \subset Y$ of polynomials of degree d , the measure μ being the (counting) probability measure on X . On V_x , the density ν_x is (of course) also the counting probability measure.

Poonen's particular application goes a little bit beyond the sieve problems described in the previous chapters: his goal is to find the density of those $f \in X$ which satisfy the sieve condition that the linear part of $\rho_x(f)$ is non-zero for each of the (usually infinitely many) closed points $x \in \Lambda$, not merely for a finite subset of them. However, the conditions required are those which classically correspond to a 'sieve of dimension 0', which means that the density $\nu_x(\Omega_x)$ of the excluded subsets tends to 0 as the degree $\deg(x)$ of x goes to infinity, sufficiently fast for the product

$$\prod_{x \in \Lambda} (1 - \nu_x(\Omega_x))$$

to converge. (The most classical example of such a situation is that of counting squarefree integers $d \geq 1$ by stating they are those which are not congruent to 0 modulo p^2 for *any* prime p , so $\{0\} \subset (\mathbf{Z}/p\mathbf{Z})^2$ corresponds to Ω_p , and has density p^{-2} .) In fact, as in this last example, Poonen is able to show in his application that the density of the (infinitely) sifted set has a limit as the degree d of the polynomials defining $X = X_d$ goes to infinity, in complete analogy with the fact that the proportion of squarefree integers among those $n \leq x$ goes to $6/\pi^2$ as $x \rightarrow +\infty$.

We conclude by stating that it is obviously possible to set up other similar sieves using other subspaces than Y_x to define the sieve setting (e.g., higher-order Taylor expansions).

Example 4.12 There are a few examples of the use of simple sieve methods in combinatorics, for instance in a paper of Liu and Murty [89] which explores a simple form of the dual sieve with some interesting combinatorial applications.

Their sieve setting amounts to taking $\Psi = (A, B, \mathbf{1}_b)$ where A and B are finite sets, and for each $b \in B$, we have a map $\mathbf{1}_b : A \rightarrow \{0, 1\}$ (in [89] the authors see (A, B) as a bipartite graph, and $\mathbf{1}_b(a) = 1$ if and only if there is an edge from a to b); the siftable set is A with identity map and counting measure, and the density is determined by $v_b(1) = |\mathbf{1}_b^{-1}(1)|/|A|$. In other words, this is also a special case of the sieve of Section 4.1, and Theorem 1 and Corollary 1 of [89] can also be trivially deduced from this (though they are simple enough to be better considered separately).

5

Degrees of representations of finite groups

5.1 Introduction

This chapter is essentially independent from the rest of the book. Indeed, it might have been another Appendix, the main difference with the appendices being that it contains mostly new results. Precisely, it is devoted to proving some inequalities which are useful in estimating quantities such as (3.13) or $R(X; \mathcal{L})$ in (3.12) when considering a group sieve (or a coset sieve) involving non-abelian finite groups G_ℓ . Indeed, we will use them later for this purpose in Chapter 7 and Chapter 8.

The reader may wish to read this introductory section only, coming back leisurely for the other parts, which can be thought of as providing a simple motivated introduction to the beautiful theory of Deligne–Lusztig characters of matrix groups over finite fields.

For motivation, consider a group sieve $(G, \Lambda, (\rho_\ell))$. Clearly, bounding the individual exponential sums $W(\pi, \tau)$ is very likely to involve the order of the groups G_m, G_n , and the degrees of their representations, which are the most basic invariants measuring the complexity of irreducible representations of a finite group, e.g., we may well obtain

$$|W(\pi, \tau)| \leq \delta(\pi, \tau)|X| + C(\dim[\pi, \bar{\tau}])^{A_1}|G_{[m,n]}|^{A_2}$$

for some constants C, A_1, A_2 (compare (7.7), Proposition 8.8). Combining those in (3.13) will then involve sums of powers of the degrees of irreducible representations of the finite groups involved. For instance, in the next chapters, we will need to bound

$$\begin{aligned} & \max_{m,\pi} \left\{ (\dim \pi) \sum_n |G_{[m,n]}| \sum_{\tau \in \Pi_n^*} (\dim \tau) \right\}, \\ & \max_{m,\pi} \left\{ (\dim \pi) \sum_n \sum_{\tau \in \Pi_n^*} (\dim \tau) \right\}. \end{aligned}$$

This motivates the following definition.

Definition 5.1 *Let G be a finite group, let $p \in [1, +\infty]$. We define*

$$A_p(G) = \left(\sum_{\rho} \dim(\rho)^p \right)^{1/p}, \quad \text{if } p \neq +\infty, \quad A_{\infty}(G) = \max_{\rho} \{\dim(\rho)\}$$

where, as everywhere in this chapter, ρ runs over all irreducible linear representations of G up to isomorphism.

We can expect to reduce our estimates to the problem of bounding $A_p(G_{[m,n]})$ in terms of m and n , for some fixed p . Indeed, we are primarily interested in $A_1(G)$ and $A_{\infty}(G)$, but $A_{5/2}(G)$ will also occur in the proof of Theorem 7.12, and other cases may be useful in other sieve settings (or for other purposes).

It is easy to give simple ‘trivial’ estimates in terms of the order of the groups themselves, which are likely to be well understood in any sieve situation. We first state these, noting that for many purposes they are certainly fine enough (this is what was used in [80]).

Proposition 5.2

- (1) *If G is abelian, we have $A_p(G) = |G|^{1/p}$ for all $p \geq 1$.*
- (2) *For any finite group G , we have $A_2(G) = |G|^{1/2}$.*
- (3) *For any finite group G , we have*

$$A_p(G) \leq |G^{\sharp}|^{1/p} A_{\infty}(G) \leq |G^{\sharp}|^{1/p} |G|^{1/2} \leq |G|^{1/2+1/p},$$

with the convention that $1/\infty = 0$, and

$$\lim_{p \rightarrow +\infty} A_p(G) = A_{\infty}(G).$$

- (4) *For any finite groups G_1 and G_2 and $p \in [1, +\infty]$, we have*

$$A_p(G_1 \times G_2) = A_p(G_1) A_p(G_2).$$

Proof

- (1) is clear since all irreducible representations of an abelian group are of dimension 1.
- (2) is simply the expression of the relation

$$A_2(G)^2 = \sum_{\rho} (\dim \rho)^2 = |G|$$

(which can be thought of, for instance, as the case $y = 1$ of (2.6) for the basis of characters of the space of conjugacy-invariant functions on G).

The first part of (3) is obtained by bounding each term in the sum defining $A_p(G)$ by the maximal value $A_\infty(G)$, and using the fact that there are as many irreducible representations as conjugacy classes; then by (2) and positivity, we have

$$A_\infty(G)^2 \leq A_2(G)^2 = |G|.$$

The limit is the standard fact that in a finite-dimensional vector space, the L^p norms converge to the L^∞ norm as $p \rightarrow +\infty$.

Finally, (4) is immediate from the description of irreducible representations of $G_1 \times G_2$ as external tensor products $\rho_1 \boxtimes \rho_2$ of irreducible representations of G_1 and G_2 respectively. □

In the next sections we will try to improve on these estimates in some cases which occur naturally in our applications. Precisely, because of (4) we need only consider the groups G_ℓ , and these are often (essentially) classical linear groups over \mathbf{F}_ℓ , such as $SL(n, \mathbf{F}_\ell)$ or symplectic groups, etc. Finding, as explicitly as possible, the irreducible representations of such groups is an important part of representation theory, where the names of Frobenius, Schur, Green, Steinberg, Deligne, and Lusztig in particular, are among the most prominent. In the remainder of this chapter, we explain what we have understood (far less than what is known!) to obtain fairly strong results concerning $A_p(G)$ for some groups of this type.

In the next sections, although we try to give concrete illustrations of all statements, which are understandable with the most basic knowledge of group theory and finite fields, it has seemed impossible to write down the arguments without employing some of the language of the theory of linear algebraic groups. For completeness, we summarize the necessary definitions in Appendix E, and we hope the concrete examples will explain clearly the results for those readers not familiar with them, and further that this may motivate them to go deeper into this beautiful theory.

5.2 Groups of Lie type with connected centres

In dealing with group sieves where G_ℓ is a finite group of Lie type, experience shows that it may not be possible to specify them exactly (in some cases, we only know that they have bounded index in $GL(n, \mathbf{F}_\ell)$ as ℓ varies, and contain $SL(n, \mathbf{F}_\ell)$, for instance; see [80] and Chapter 8). Our results are biased to this

case, and we start with an easy monotonicity lemma that is helpful to deal with such discrepancies.

Lemma 5.3 *Let G be a finite group and $H \subset G$ a subgroup, $p \in [1, +\infty]$. We have*

$$A_p(H) \leq A_p(G).$$

Proof For any irreducible representation ρ of H , choose (arbitrarily) an irreducible representation π_ρ of G that occurs with positive multiplicity in the induced representation $\text{Ind}_H^G \rho$.

Let π be a representation of G of the form π_{ρ_0} for some representation ρ_0 . For any ρ where $\pi_\rho = \pi$, we have

$$\left\langle \rho, \text{Res}_H^G \pi \right\rangle_H = \left\langle \text{Ind}_H^G \rho, \pi \right\rangle_G > 0,$$

by Frobenius reciprocity, i.e., all ρ with $\pi_\rho = \pi$ occur in the restriction of π to H , so that $\dim \pi$ is at least as large as the sum of the $\dim \rho$ over ρ with $\pi_\rho = \pi$. More generally, for $p \neq +\infty$, we obtain

$$\sum_{\pi_\rho=\pi} \dim(\rho)^p \leq \left(\sum_{\pi_\rho=\pi} \dim(\rho) \right)^p \leq \dim(\pi)^p,$$

and summing over all representations of the form π_ρ gives the inequality

$$A_p(H)^p \leq A_p(G)^p$$

by positivity. This settles the case $p \neq +\infty$, and the other case only requires noticing that $\dim \rho \leq \dim \pi_\rho \leq A_\infty(G)$, since ρ occurs in the restriction of π_ρ . \square

We come to the main result of this chapter. The terminology, which may not be familiar to all readers, is explained by examples after the proof, and reviewed quickly in Appendix E. There should be no confusion between p as used earlier and the characteristic of the finite field \mathbf{F}_q which occurs here, since from now on we will mostly work with A_1 and A_∞ .

Proposition 5.4 *Let \mathbf{G}/\mathbf{F}_q be a split connected reductive linear algebraic group of dimension d and rank r over a finite field, with connected centre. Let W be its Weyl group and $G = \mathbf{G}(\mathbf{F}_q)$ the finite group of rational points of \mathbf{G} .*

(1) *For any subgroup $H \subset G$ and $p \in [1, +\infty]$, we have*

$$A_p(H) \leq (q+1)^{(d-r)/2+r/p} \left(1 + \frac{2r|W|}{q-1} \right)^{1/p},$$

with the convention $r/p = 0$ if $p = +\infty$, in particular the second factor is equal to 1 for $p = +\infty$.

- (2) If \mathbf{G} is a product of groups of type A or C , i.e., of linear and symplectic groups, then

$$A_p(H) \leq (q+1)^{(d-r)/2+r/p}.$$

The proof is based on a simple interpolation argument from the extreme cases $p = 1$, $p = +\infty$. Indeed by Lemma 5.3 we can clearly assume $H = G$ and by writing the obvious inequality

$$A_p(G)^p = \sum_{\rho} \dim(\rho)^p \leq A_{\infty}(G)^{p-1} A_1(G),$$

we see that it suffices to prove the following:

Proposition 5.5 *Let \mathbf{G}/\mathbf{F}_q be a split connected reductive linear algebraic group of dimension d with connected centre, and let $G = \mathbf{G}(\mathbf{F}_q)$ be the finite group of its rational points. Let r be the rank of \mathbf{G} . Then we have*

$$A_{\infty}(G) \leq \frac{|G|_{p'}}{(q-1)^r} \leq (q+1)^{(d-r)/2}, \quad A_1(G) \leq (q+1)^{(d+r)/2} \left(1 + \frac{2r|W|}{q-1}\right), \quad (5.1)$$

where $n_{p'}$ denotes the prime-to- p part of a rational number n , p being the characteristic of \mathbf{F}_q . Moreover, if the principal series of G is not empty,¹ there is equality, so that

$$A_{\infty}(G) = \frac{|G|_{p'}}{(q-1)^r}$$

and $\dim \rho = A_{\infty}(G)$, if and only if ρ is in the principal series.

Finally if \mathbf{G} is a product of groups of type A or C , then the second factor $1 + 2r|W|/(q-1)$ may be removed in the bound for $A_1(G)$.

It seems very possible that the factor $1 + 2r|W|/(q-1)$ could always be removed, but we haven't been able to figure this out using Deligne–Lusztig characters, and in fact for groups of type A or C , we simply quote *exact formulas* for $A_1(G)$ due to Gow, Klyachko and Vinroot, which are proved in completely different ways. The extra factor is not likely to be a problem in many applications where $q \rightarrow +\infty$, but it may be questionable for uniformity with respect to the rank, because $|W|$ typically grows super-exponentially with r .

The ideas in the proof were suggested and explained by J. Michel.

¹ In particular if q is large enough given \mathbf{G} .

Proof This is based on properties of the Deligne–Lusztig generalized characters. We will mostly refer to [31] and [19] for all facts which are needed (using notation from [31], except for writing simply G instead of \mathbf{G}^F as used there). We identify irreducible representations of G (up to isomorphism) with their characters, seen as complex-valued class functions on G , i.e., conjugacy-invariant functions on G .

First, for a connected reductive group \mathbf{G}/\mathbf{F}_q over a finite field, Deligne and Lusztig have constructed (see, e.g., [31, 11.14]) a family $R_{\mathbf{T}}^{\mathbf{G}}(\theta)$ of generalized representations of $G = \mathbf{G}(\mathbf{F}_q)$ (i.e., linear combinations with integer coefficients of ‘genuine’ representations of G), parametrized by pairs (\mathbf{T}, θ) consisting of a maximal torus $\mathbf{T} \subset \mathbf{G}$ defined over \mathbf{F}_q and a (one-dimensional) character θ of the finite abelian group $T = \mathbf{T}(\mathbf{F}_q)$. The $R_{\mathbf{T}}^{\mathbf{G}}(\theta)$ are not all irreducible, but any irreducible character occurs (with positive or negative multiplicity) in the decomposition of at least one such generalized character. Moreover, $R_{\mathbf{T}}^{\mathbf{G}}(\theta)$ only depends (as a class function) on the G -conjugacy class of the pair (\mathbf{T}, θ) .

We quote here a useful classical fact: for any \mathbf{T} we have

$$(q-1)^r \leq |T| \leq (q+1)^r \quad (5.2)$$

(see, e.g. [31, 13.7 (ii)]), and moreover $|T| = (q-1)^r$ if and only if \mathbf{T} is a split torus (i.e., $\mathbf{T} \simeq \mathbf{G}_m^r$ over \mathbf{F}_q). Indeed, we have

$$|T| = |\det(q^n - w \mid Y_0)|$$

where $w \in W$ is such that \mathbf{T} is obtained from a split torus \mathbf{T}_0 by ‘twisting with w ’ (see, e.g. [19, Proposition 3.3.5]), and $Y_0 \simeq \mathbf{Z}^r$ is the group of cocharacters of \mathbf{T}_0 . If $\lambda_1, \dots, \lambda_r$ are the eigenvalues of w acting on Y_0 , which are roots of unity, then we have

$$|T| = \prod_{i=1}^r (q - \lambda_i),$$

and so $|T| = (q-1)^r$ if and only if each λ_i is equal to 1, if and only if w acts trivially on Y_0 , if and only if $w = 1$ (W acts faithfully on Y_0) and \mathbf{T} is split.

As in [31, 12.12], we denote by $\rho \mapsto p(\rho)$ the orthogonal projection of the space $L^2(G^{\#}) \subset L^2(G)$ of complex-valued conjugacy-invariant functions on G to the subspace generated by Deligne–Lusztig characters, where $L^2(G^{\#})$ is given the standard inner product

$$\langle f, g \rangle = \frac{1}{|G|} \sum_{x \in G} f(x) \overline{g(x)},$$

(already seen in Chapter 3) and for a representation ρ , we of course denote by $p(\rho) = p(\text{Tr } \rho)$ the projection of its character.

For any representation ρ , we have $\dim(\rho) = \dim(p(\rho))$, where $\dim(f)$, for an arbitrary function $f \in L^2(G^\sharp)$ is obtained by linearity from the degree of characters. Indeed, for any f , standard character theory shows that

$$\dim(f) = \langle f, \text{reg}_G \rangle$$

where reg_G is the regular representation of G . The regular representation is in the subspace spanned by the Deligne–Lusztig characters (see, e.g., [31, 12.14]), so by definition of an orthogonal projector we have

$$\dim(\rho) = \langle \rho, \text{reg}_G \rangle = \langle p(\rho), \text{reg}_G \rangle = \dim(p(\rho)).$$

Now because the characters $R_{\mathbf{T}}^G(\theta)$ for distinct conjugacy classes of (\mathbf{T}, θ) are orthogonal (see, e.g. [31, 11.15]), we can write

$$p(\rho) = \sum_{(\mathbf{T}, \theta)} \frac{\langle \rho, R_{\mathbf{T}}^G(\theta) \rangle}{\langle R_{\mathbf{T}}^G(\theta), R_{\mathbf{T}}^G(\theta) \rangle} R_{\mathbf{T}}^G(\theta)$$

(sum over all distinct Deligne–Lusztig characters) and so

$$\dim(p(\rho)) = \sum_{(\mathbf{T}, \theta)} \frac{\langle \rho, R_{\mathbf{T}}^G(\theta) \rangle}{\langle R_{\mathbf{T}}^G(\theta), R_{\mathbf{T}}^G(\theta) \rangle} \dim(R_{\mathbf{T}}^G(\theta)).$$

By [31, 12.9] we have

$$\dim(R_{\mathbf{T}}^G(\theta)) = \varepsilon_G \varepsilon_{\mathbf{T}} |G|_{p'} |T|^{-1}, \quad (5.3)$$

where $\varepsilon_G = (-1)^r$ and $\varepsilon_{\mathbf{T}} = (-1)^{r(\mathbf{T})}$, $r(\mathbf{T})$ being the \mathbf{F}_q -rank of \mathbf{T} (see [31, p. 65] or Appendix E for the definition). This yields the formula

$$\dim(p(\rho)) = |G|_{p'} \sum_{(\mathbf{T}, \theta)} \frac{1}{|T|} \frac{\langle \rho, \varepsilon_G \varepsilon_{\mathbf{T}} R_{\mathbf{T}}^G(\theta) \rangle}{\langle R_{\mathbf{T}}^G(\theta), R_{\mathbf{T}}^G(\theta) \rangle}. \quad (5.4)$$

Now we use the fact that pairs (\mathbf{T}, θ) are partitioned into *geometric conjugacy classes*, defined as follows: two pairs (\mathbf{T}, θ) and (\mathbf{T}', θ') are geometrically conjugate if and only if there exists $g \in \mathbf{G}(\bar{\mathbf{F}}_q)$ such that $\mathbf{T} = g\mathbf{T}'g^{-1}$ and for all n such that $g \in \mathbf{G}(\mathbf{F}_{q^n})$, we have

$$\theta(N_{\mathbf{F}_{q^n}/\mathbf{F}_q}(x)) = \theta'(N_{\mathbf{F}_{q^n}/\mathbf{F}_q}(g^{-1}xg)) \quad \text{for } x \in \mathbf{T}(\mathbf{F}_{q^n})$$

(see, e.g. [31, 13.2]). The point is the following property of geometric conjugacy classes: if the generalized characters $R_{\mathbf{T}}^G(\theta)$ and $R_{\mathbf{T}'}^G(\theta')$ have a common irreducible component, then (\mathbf{T}, θ) and (\mathbf{T}', θ') are geometrically conjugate (see, e.g. [31, 13.2]).

In particular, for a given irreducible representation ρ , if $\langle \rho, R_{\mathbf{T}}^G(\theta) \rangle$ is non-zero for some (\mathbf{T}, θ) , then only pairs (\mathbf{T}', θ') geometrically conjugate to (\mathbf{T}, θ) may satisfy $\langle \rho, R_{\mathbf{T}'}^G(\theta) \rangle \neq 0$. So we have

$$\dim(p(\rho)) = |G|_{p'} \sum_{(\mathbf{T}, \theta) \in \kappa} \frac{1}{|T|} \frac{\langle \rho, \varepsilon_G \varepsilon_{\mathbf{T}} R_{\mathbf{T}}^G(\theta) \rangle}{\langle R_{\mathbf{T}}^G(\theta), R_{\mathbf{T}}^G(\theta) \rangle},$$

for some geometric conjugacy class κ , depending on ρ . By Cauchy–Schwarz, we obtain

$$\dim(p(\rho)) \leq |G|_{p'} \left(\sum_{(\mathbf{T}, \theta) \in \kappa} \frac{1}{|T|^2} \frac{1}{\langle R_{\mathbf{T}}^G(\theta), R_{\mathbf{T}}^G(\theta) \rangle} \right)^{1/2} \left(\sum_{(\mathbf{T}, \theta) \in \kappa} \frac{|\langle \rho, R_{\mathbf{T}}^G(\theta) \rangle|^2}{\langle R_{\mathbf{T}}^G(\theta), R_{\mathbf{T}}^G(\theta) \rangle} \right)^{1/2}.$$

The second term on the right is simply $\langle p(\rho), p(\rho) \rangle \leq \langle \rho, \rho \rangle = 1$. As for the first term we have

$$\sum_{(\mathbf{T}, \theta) \in \kappa} \frac{1}{|T|^2} \frac{1}{\langle R_{\mathbf{T}}^G(\theta), R_{\mathbf{T}}^G(\theta) \rangle} \leq \frac{1}{(q-1)^{2r}} \sum_{(\mathbf{T}, \theta) \in \kappa} \frac{1}{\langle R_{\mathbf{T}}^G(\theta), R_{\mathbf{T}}^G(\theta) \rangle}$$

by (5.2). Now it is known that for each class κ , the assumption that \mathbf{G} has connected centre implies that the generalized character

$$\chi(\kappa) = \sum_{(\mathbf{T}, \theta) \in \kappa} \frac{\varepsilon_G \varepsilon_{\mathbf{T}} R_{\mathbf{T}}^G(\theta)}{\langle R_{\mathbf{T}}^G(\theta), R_{\mathbf{T}}^G(\theta) \rangle}$$

is in fact an irreducible character of G (such characters are called *regular* characters;² see, e.g., [19, Proposition 8.4.7]). This implies that

$$\sum_{(\mathbf{T}, \theta) \in \kappa} \frac{1}{\langle R_{\mathbf{T}}^G(\theta), R_{\mathbf{T}}^G(\theta) \rangle} = \langle \chi(\kappa), \chi(\kappa) \rangle = 1,$$

and so we have

$$\dim p(\rho) \leq \frac{|G|_{p'}}{(q-1)^r}. \quad (5.5)$$

Now observe that we will have equality in this argument if ρ is itself of the form $\pm R_{\mathbf{T}}^G(\theta)$, and if $|T| = (q-1)^r$. These conditions hold for representations of the principal series, i.e., characters $R_{\mathbf{T}}^G(\theta)$ for an \mathbf{F}_q -split torus \mathbf{T} and a character θ ‘in general position’ (see, e.g., [19, Corollary 7.3.5]). Such characters are also, more elementarily, induced characters $\text{Ind}_B^G(\theta)$, where $B = \mathbf{B}(\mathbf{F}_q)$ is a Borel subgroup containing T , for some Borel subgroup \mathbf{B} defined over \mathbf{F}_q containing \mathbf{T} (which exist for a split torus \mathbf{T}) and θ is extended to B by setting $\theta(u) = 1$ for unipotent elements $u \in B$. For this, see, e.g., [94, Proposition 2.6].

² Not to be confused with ‘the’ regular representation reg_G .

Conversely, let ρ be such that

$$\dim \rho = \frac{|G|_{p'}}{(q-1)^r}$$

and let κ be the associated geometric conjugacy class. From the above, for any (\mathbf{T}, θ) in κ , we have $|T| = (q-1)^r$, i.e., \mathbf{T} is \mathbf{F}_q -split. Now it follows from Lemma 5.6 below (probably well known) that this implies that $R_{\mathbf{T}}^{\mathbf{G}}(\theta)$ is an irreducible representation, so must be equal to ρ .

We now come to $A_1(G)$. To deal with the fact that, in (5.4), $|T|$ depends on $(\mathbf{T}, \theta) \in \kappa$, we write

$$\begin{aligned} \dim(p(\rho)) &= \frac{|G|_{p'}}{(q-1)^r} \sum_{\kappa} \langle \rho, \chi(\kappa) \rangle \\ &\quad + |G|_{p'} \sum_{(\mathbf{T}, \theta)} \left(\frac{1}{|T|} - \frac{1}{(q-1)^r} \right) \frac{\varepsilon_{\mathbf{G}} \varepsilon_{\mathbf{T}}(\rho, R_{\mathbf{T}}^{\mathbf{G}}(\theta))}{\langle R_{\mathbf{T}}^{\mathbf{G}}(\theta), R_{\mathbf{T}}^{\mathbf{G}}(\theta) \rangle} \end{aligned} \quad (5.6)$$

(since by (5.2), the dependency is rather weak).

Now summing over ρ , consider the first term's contribution. Since $\chi(\kappa)$ is an irreducible character, the sum

$$\sum_{\rho} \sum_{\kappa} \langle \rho, \chi(\kappa) \rangle$$

is simply the number of geometric conjugacy classes. This is given by $q^{r'}|Z|$ by [31, 14.42] or [19, Theorem 4.4.6 (ii)], where r' is the semisimple rank of \mathbf{G} and $Z = Z(\mathbf{G})(\mathbf{F}_q)$ is the group of rational points of the centre of \mathbf{G} . For this quantity, note that the centre of \mathbf{G} being connected implies that $Z(\mathbf{G})$ is the radical of \mathbf{G} (see, e.g., [125, Proposition 7.3.1]) so $Z(\mathbf{G})$ is a torus and $r = r' + \dim Z(\mathbf{G})$. So using again the bounds (5.2) for the cardinality of the group of rational points of a torus, we obtain

$$|Z|q^{r'} \leq (q+1)^r. \quad (5.7)$$

To estimate the sum of the contributions in the second term, say $\sum t(\rho)$, we write

$$\sum_{\rho} t(\rho) = |G|_{p'} \sum_{(\mathbf{T}, \theta)} \left(\frac{1}{|T|} - \frac{1}{(q-1)^r} \right) \frac{\varepsilon_{\mathbf{G}} \varepsilon_{\mathbf{T}}(\sum_{\rho} \rho, R_{\mathbf{T}}^{\mathbf{G}}(\theta))}{\langle R_{\mathbf{T}}^{\mathbf{G}}(\theta), R_{\mathbf{T}}^{\mathbf{G}}(\theta) \rangle},$$

and we bound

$$\left| \left\langle \sum_{\rho} \rho, R_{\mathbf{T}}^{\mathbf{G}}(\theta) \right\rangle \right| \leq \langle R_{\mathbf{T}}^{\mathbf{G}}(\theta), R_{\mathbf{T}}^{\mathbf{G}}(\theta) \rangle \quad (5.8)$$

for any (\mathbf{T}, θ) , since we can write

$$R_{\mathbf{T}}^{\mathbf{G}}(\theta) = \sum_{\rho} a(\rho) \rho \quad \text{with } a(\rho) \in \mathbf{Z},$$

and therefore

$$\left| \left\langle \sum_{\rho} \rho, R_{\mathbf{T}}^{\mathbf{G}}(\theta) \right\rangle \right| = \left| \sum_{\rho} a(\rho) \right| \leq \sum_{\rho} |a(\rho)|^2 = \langle R_{\mathbf{T}}^{\mathbf{G}}(\theta), R_{\mathbf{T}}^{\mathbf{G}}(\theta) \rangle. \quad (5.9)$$

Thus

$$\sum_{\rho} t(\rho) \leq \frac{|G|_{p'}}{(q-1)^r} \frac{2r}{q-1} |\{(\mathbf{T}, \theta)\}|.$$

There are at most $|W|$ different choices of \mathbf{T} up to G -conjugacy, and for each there are at most $|T| \leq (q+1)^r$ different characters, and so we have

$$\sum_{\rho} t(\rho) \leq \frac{|G|_{p'}}{(q-1)^r} \frac{2r|W|}{q-1} (q+1)^r, \quad (5.10)$$

and

$$\sum_{\rho} \dim \rho \leq (q+1)^r \frac{|G|_{p'}}{(q-1)^r} \left(1 + \frac{2r|W|}{q-1}\right). \quad (5.11)$$

To conclude, we use the classical formula

$$|G| = q^N \prod_{1 \leq i \leq r} (q^{d_i} - 1),$$

where N is the number of positive roots of \mathbf{G} , and the d_i are the degrees of invariants of the Weyl group (this is because \mathbf{G} is split; see, e.g. [19, 2.4.1 (iv); 2.9, p. 75]). So

$$|G|_{p'} = \prod_{1 \leq i \leq r} (q^{d_i} - 1)$$

and

$$\begin{aligned} \frac{|G|_{p'}}{(q-1)^r} &= \prod_{1 \leq i \leq r} \frac{q^{d_i} - 1}{q-1} \leq \prod_{1 \leq i \leq r} (q+1)^{d_i-1} \\ &= (q+1)^{\sum (d_i-1)} = (q+1)^{(d-r)/2}, \end{aligned} \quad (5.12)$$

since $\sum (d_i - 1) = N$ and $N = (d - r)/2$ (see, e.g. [19, 2.4.1], [125, 8.1.3]).

Inserting this in (5.5) we derive the first inequality in (5.1), and with (5.11), we get

$$A_1(G) \leq (q+1)^{(d+r)/2} \left(1 + \frac{2r|W|}{q-1}\right),$$

which is the second part of (5.1).

Now we explain why the extra factor involving the Weyl group can be removed for products of groups of type A and C . Clearly it suffices to work with $\mathbf{G} = GL(n)$ and $\mathbf{G} = CSp(2g)$.

For $\mathbf{G} = GL(n)$, with $d = n^2$ and $r = n$, Gow [50] and Klyachko [78] have proved independently that $A_1(G)$ is equal to the number of symmetric matrices in G . The bound

$$A_1(G) \leq (q+1)^{(n^2+n)/2}$$

follows immediately.

For $\mathbf{G} = CSp(2g)$, with $d = 2g^2 + g + 1$ and $r = g + 1$, the exact analogue of Gow's theorem is due to Vinroot [129]. Again, Vinroot's result implies $A_1(G) \leq (q+1)^{(d+r)/2}$ in this case (see [129, Corollary 6.1], and use the formulas for the order of unitary and linear groups to check the final bound). \square

Here is the lemma used in the determination of $A_\infty(G)$ when there is a character in general position of a split torus:

Lemma 5.6 *Let \mathbf{G}/\mathbf{F}_q be a split connected reductive linear algebraic group of dimension d and let $G = \mathbf{G}(\mathbf{F}_q)$ be the finite group of its rational points. Let \mathbf{T} be a split torus in \mathbf{G} , θ a character of $T = \mathbf{T}(\mathbf{F}_q)$. If \mathbf{T}' is also a split torus for any pair (\mathbf{T}', θ') geometrically conjugate to (\mathbf{T}, θ) , then $R_{\mathbf{T}}^G(\theta)$ is irreducible.*

Proof If $R_{\mathbf{T}}^G(\theta)$ is not irreducible, then by the inner product formula for Deligne–Lusztig characters, there exists $w \in W$, $w \neq 1$, such that ${}^w\theta = \theta$ (using the natural action of W on the characters of \mathbf{T} ; see, e.g. [31, Corollary 11.15]). Let \mathbf{T}' be a torus obtained from \mathbf{T} by ‘twisting by w ’, i.e., $\mathbf{T}' = g\mathbf{T}g^{-1}$ where $g \in \mathbf{G}$ is such that $g^{-1}\text{Fr}(g) = w$ (see, e.g. [19, 3.3]). Let $Y = \text{Hom}(\mathbf{G}_m, \mathbf{T}) \simeq \mathbf{Z}'$ (respectively Y') be the abelian group of cocharacters of \mathbf{T} (respectively \mathbf{T}'); the conjugation isomorphism $\mathbf{T} \rightarrow \mathbf{T}'$ gives rise to a conjugation isomorphism $Y \rightarrow Y'$ ([19]). Moreover, there is an action of the Frobenius automorphism Fr on Y and a canonical isomorphism $T \simeq Y/(\text{Fr} - 1)Y$ (see, e.g. [31, Proposition 13.7]), hence canonical isomorphisms of the character groups \hat{T} and \hat{T}' as subgroups of the characters groups of Y and Y' :

$$\hat{T} \simeq \{\chi : Y \rightarrow \mathbf{C}^\times \mid (\text{Fr} - 1)Y \subset \text{Ker } \chi\},$$

$$\hat{T}' \simeq \{\chi : Y' \rightarrow \mathbf{C}^\times \mid (\text{Fr} - 1)Y' \subset \text{Ker } \chi\}.$$

Unraveling the definitions, a simple calculation shows that the condition ${}^w\theta = \theta$ is precisely what is needed to prove that the character χ of Y associated

to θ , when ‘transported’ to a character χ' of Y' by the conjugation isomorphism, still satisfies $\text{Ker } \chi' \supset (\text{Fr} - 1)Y'$ (see in particular [19, Proposition 3.3.4]), so is associated with a character $\theta' \in \hat{T}'$.

Using the characterization of geometric conjugacy in [31, Proposition 13.8], it is then clear that (\mathbf{T}, θ) is geometrically conjugate to (\mathbf{T}', θ') , and since $w \neq 1$, the torus \mathbf{T}' is not split. So by contraposition, the lemma is proved. \square

Remark 5.7 Characters of a split torus \mathbf{T} in general position can only exist if $|T| > r$ since it is necessary that T has r distinct characters. One may therefore wonder what happens for fixed q if \mathbf{G} runs over a family with $r \rightarrow +\infty$, e.g., for $GL(r, \mathbf{F}_q)$. At the very least, we have

$$A_\infty(\mathbf{G}(\mathbf{F}_q)) \geq q^{(d-r)/2},$$

for any reductive group \mathbf{G}/\mathbf{F}_q (split or not), because of the existence of the important *Steinberg character* St_G . Indeed, this character is always defined for a reductive group \mathbf{G} over a finite field, and is an irreducible character of degree equal to the order of a p -Sylow subgroup of $G = \mathbf{G}(\mathbf{F}_q)$ (see, e.g., [19, Corollary 6.3], [31, Corollary 9.3]), namely

$$\dim \text{St}_G = q^{(d-r)/2}.$$

In terms of Deligne–Lusztig characters, St_G is a component of $R_{\mathbf{T}}^G(1)$ for any maximal torus \mathbf{T} . (Note that, in contrast to many irreducible characters, which are only known as class functions on the group G , the Steinberg *representation*, namely, an actual vector space on which G acts according to St_G , can be described in fairly explicit terms.)

Note that, going rather in the opposite direction of what we have discussed, an important question in the representation theory of finite groups of Lie type is to find a *lower bound* for the minimal dimension of an irreducible representation (which is not of dimension 1; if $\mathbf{G} = SL(r)$, this amounts to asking for the minimal dimension of a non-trivial irreducible representation). We will not make use of such information in this book, but this turns out to be important, for instance, in some arguments used to prove the existence of a ‘spectral gap’ in certain families of graphs or hyperbolic surfaces (e.g., to prove a form of Selberg’s theorem on the smallest eigenvalue of a congruence quotient $\Gamma(p) \backslash \mathbf{H}$, the idea being that ‘exceptional eigenvalues’ must have high multiplicity because the covering group $SL(2, \mathbf{F}_p) = \Gamma(1)/\Gamma(p)$ acts without invariant vectors on the corresponding Laplace eigenspace;³ see for

³ For $SL(2, \mathbf{F}_q)$, the character table shows that a non-trivial irreducible representation has degree $\geq \frac{1}{2}(q-1)$.

instance [114] and [47] for such arguments). We will see in Chapter 7 that families of graphs with a spectral gap are another fertile source of applications of the large sieve.

5.3 Examples

Here are the basic examples of reductive groups with connected centres that we will use.

Example 5.8

- (1) Let ℓ be prime, $r \geq 1$ and let $\mathbf{G} = GL(r)/\mathbf{F}_\ell$. Then $G = GL(r, \mathbf{F}_\ell)$, \mathbf{G} is a split connected reductive group of rank r , dimension r^2 , with connected centre of dimension 1. So from Lemma 5.3 and Proposition 5.4, we get

$$A_p(H) \leq (\ell + 1)^{r(r-1)/2 + r/p}$$

for $p \in [1, +\infty]$ for any subgroup H of G , and in particular

$$A_\infty(H) \leq (\ell + 1)^{r(r-1)/2} \quad \text{and} \quad A_1(H) \leq (\ell + 1)^{r(r+1)/2}.$$

It would be interesting to know if there are other values of p besides $p = 1, 2$ and $+\infty$ (the latter when q is large enough) for which $A_p(GL(n, \mathbf{F}_q))$ can be computed exactly.

- (2) Let $\ell \neq 2$ be prime, $g \geq 1$ and let $\mathbf{G} = CSp(2g)/\mathbf{F}_\ell$. Then $G = CSp(2g, \mathbf{F}_\ell)$ and G is a split connected reductive group of rank $g + 1$ and dimension $2g^2 + g + 1$, with connected centre. So from Lemma 5.3 and Proposition 5.4, we get

$$A_p(H) \leq (\ell + 1)^{g^2 + (g+1)/p}$$

for $p \in [1, +\infty]$ for any subgroup H of G , and in particular

$$A_\infty(H) \leq (\ell + 1)^{g^2} \quad \text{and} \quad A_1(H) \leq (\ell + 1)^{g^2 + g + 1}.$$

- (3) For some ‘small rank’ groups, the character tables are completely known, and therefore the exact computation of A_p is possible for all p (even for complex p , if desired). For example, in the case of $GL(2, \mathbf{F}_q)$ over fields of odd characteristic, we have

$$\sum_{\rho} \dim(\rho)^p = (q-1)(q^p+1) + \frac{(q-1)(q-2)}{2}(q+1)^p + \frac{q(q-1)}{2}(q-1)^p,$$

for all $p \in \mathbf{C}$ (see Section C.4 for the character table of $GL(2, \mathbf{F}_q)$).

Computing exactly A_∞ and A_1 for $\mathbf{G} = GL(2)$ and $\mathbf{G} = SL(2)$ (see the character tables of $GL(2, \mathbf{F}_q)$ and $SL(2, \mathbf{F}_q)$ for q odd), one finds that:

$$A_\infty(GL(2, \mathbf{F}_q)) = \begin{cases} q & \text{if } q = 2, \\ q + 1 & \text{if } q > 2, \end{cases} \quad A_1(GL(2, \mathbf{F}_q)) = q^3 - q^2$$

$$A_\infty(SL(2, \mathbf{F}_q)) = \begin{cases} q & \text{if } q = 2, 3, \\ q + 1 & \text{if } q > 3, \end{cases} \quad A_1(SL(2, \mathbf{F}_q)) = q^2 + q.$$

By multiplicativity, the case of prime q implies

$$A_\infty(GL(2, \mathbf{Z}/m\mathbf{Z})) = \begin{cases} \psi(m) & \text{if } m \text{ is odd,} \\ 2\psi(m/2) & \text{if } m \text{ is even,} \end{cases} \leq \psi(m), \quad (5.13)$$

$$A_1(GL(2, \mathbf{Z}/m\mathbf{Z})) = m^2\varphi(m), \quad (5.14)$$

$$A_\infty(SL(2, \mathbf{Z}/m\mathbf{Z})) = \begin{cases} \psi(m) & \text{if } (m, 6) = 1 \\ 2\psi(m/2) & \text{if } m \text{ is even} \\ 2\psi(m/3) & \text{if } m \equiv 0 \pmod{3} \\ 4\psi(m/6) & \text{if } m \equiv 0 \pmod{6}, \end{cases} \leq \psi(m), \quad (5.15)$$

$$A_1(SL(2, \mathbf{Z}/m\mathbf{Z})) = m\psi(m), \quad (5.16)$$

for all squarefree integers $m \geq 1$ (where $\psi(m)$ is defined in the section on notation). We will use this in Section 7.4.

For $GL(3)$ and $GL(4)$, $SL(3)$ and $SL(4)$, one can look at [126]; the case of $Sp(4)$ is also fairly classical, the character table being due to Srinivasan (this is where the first example of a so-called cuspidal unipotent irreducible representation occurs; there are no such representations for $GL(n)$).

5.4 Some groups with disconnected centres

In the case of $G = SL(r, \mathbf{F}_q)$ or $G = Sp(2g, \mathbf{F}_q)$, which correspond to \mathbf{G} where the centre is not connected, the bound for $A_\infty(G)$ given by Example 5.8 is still sharp if we see G as subgroup of $GL(r, \mathbf{F}_q)$ or $CSp(2g, \mathbf{F}_q)$, because both d and r increase by 1, so $d - r$ doesn't change. However, for $A_1(G)$, the exponent increases by one. Here is a slightly different argument that almost recovers the 'right' bound.

Lemma 5.9 *Let $\mathbf{G} = SL(n)$ or $Sp(2g)$ over \mathbf{F}_q , let d be the dimension and r the rank of \mathbf{G} , and $G = \mathbf{G}(\mathbf{F}_q)$. Then we have the following bounds*

$$A_p(G) \leq \kappa^{1/p} (q + 1)^{(d-r)/2+r/p} \left(\frac{q+1}{q-1} \right)^{1/p}.$$

and

$$A_p(G) \leq (q+1)^{(d-r)/2+r/p} \left(\frac{q+1}{q-1} \right)^{1/p} \left(1 + \frac{2\kappa(r+1)|W|}{q-1} \right)^{1/p}$$

for any $p \in [1, +\infty]$, where $\kappa = n$ for $SL(n)$ and $\kappa = 2$ for $Sp(2g)$.

The first bound is better for fixed q , whereas the second is almost as sharp as the bound for $GL(n)$ or $CSp(2g)$ if q is large.

Proof As we observed before the statement, this holds for $p = +\infty$, so it suffices to consider $p = 1$ and then use the same interpolation argument as for Proposition 5.4.

Let $G_1 = GL(n)$ or $CSp(2g)$ for $G = SL(n)$ or $Sp(2g)$ respectively, $G_1 = G_1(\mathbf{F}_q)$. We use the exact sequence

$$1 \rightarrow G \rightarrow G_1 \xrightarrow{m} \Gamma = \mathbf{F}_q^\times \rightarrow 1$$

(compare with Section 3.3) where m is either the determinant of a matrix or the multiplier of a symplectic similitude. Let ρ be an irreducible representation of G , and as in the proof of Lemma 5.3, let π_ρ be any irreducible component in the induced representation $\text{Ind}_G^{G_1} \rho$. The point is that all ‘twists’ $\pi_\rho \otimes \psi$, where ψ is a character of \mathbf{F}_q^\times lifted to G_1 through m , are isomorphic restricted to G , and hence each $\pi_\rho \otimes \psi$ contains ρ when restricted to G , and contains even all ρ with the same π_ρ . So if $\pi \sim \pi'$, for representations of G_1 , denotes isomorphism when restricted to G , we have

$$A_1(G) \leq \sum_{\{\pi\}/\sim} \dim \pi$$

where the sum is over a set of representatives for this equivalence relation. On the other hand, $\dim \pi = \dim \pi'$ for $\pi \sim \pi'$, and for each π there are $|\hat{\Gamma}/\hat{\Gamma}^\pi|$ distinct representations equivalent to π , with notation as in Lemma 3.2. Hence,

$$A_1(G) \leq \frac{1}{q-1} \sum_{\pi} |\hat{\Gamma}^\pi| \dim \pi.$$

From, e.g., [80, Lemma 2.3], we know that $\hat{\Gamma}^\pi$ has order at most n (for $SL(n)$) or 2 (for $Sp(2g)$), which by applying Proposition 5.4 yields the first bound,⁴ namely

$$A_1(G) \leq \kappa \frac{(q+1)^{(d+r)/2}}{q-1}, \quad \text{with } \kappa = 2 \text{ or } n.$$

⁴ This suffices for the applications in this book.

To obtain the refined estimate, observe that in the formula (5.6) for the dimension of an irreducible representation ρ of G_1 , the first term is zero unless ρ is a regular character, and the second $t(\rho)$ is smaller by a factor of size roughly equal to q . If π is regular, we have $\hat{\Gamma}^\pi = 1$ by Lemma 5.10 below. So it follows that

$$\begin{aligned} A_1(G) &\leq \frac{1}{q-1} \left\{ \sum_{\pi \text{ regular}} \dim \pi + \kappa \sum_{\pi \text{ not regular}} \dim \pi \right\} \\ &\leq \frac{A_\infty(G_1)}{q-1} q^r (q-1) + \kappa \sum_{\pi \text{ not regular}} t(\rho) \end{aligned}$$

(in the first term, $q^r(q-1)$ is the number of geometric conjugacy classes for G_1 , computed as in (5.7), since r is the semisimple rank of G_1). We have the analogue of (5.10):

$$\begin{aligned} \sum_{\pi \text{ not regular}} t(\pi) &\leq \frac{|G_1|_{p'}}{(q-1)^{r+1}} \frac{2(r+1)|W|}{q-1} (q+1)^{r+1} \\ &\leq 2(r+1)|W| \frac{(q+1)^{(d+r)/2+1}}{q-1}, \end{aligned}$$

by (5.12) (because

$$\left| \left\langle \sum_{\pi \text{ not regular}} \pi, R_{\mathbf{T}}^{G_1}(\theta) \right\rangle \right| \leq \langle R_{\mathbf{T}}^{G_1}(\theta), R_{\mathbf{T}}^{G_1}(\theta) \rangle,$$

see (5.9), and the same argument leading to (5.8)). The bound

$$A_1(G) \leq (q+1)^{(d+r)/2} \left(1 + \frac{2\kappa(r+1)|W|}{q-1} \right)$$

follows. □

Here is the lemma, also unlikely to be very new, that we used in the proof:

Lemma 5.10 *Let $\mathbf{G} = GL(n)$ or $CSp(2g)$ over \mathbf{F}_q , $G = \mathbf{G}(\mathbf{F}_q)$. For any regular irreducible character ρ of G , we have $\hat{\Gamma}^\rho = 1$.*

Proof As above, let $m : \mathbf{G} \rightarrow \mathbf{G}_m$ be the determinant or multiplier character. Let ρ be a regular character and ψ a character of \mathbf{F}_q^\times such that $\rho \otimes \psi \simeq \rho$, where ψ is shorthand for $\psi \circ m$. We wish to show that ψ is trivial to conclude $\hat{\Gamma}^\rho = 1$. For this purpose, write

$$\rho = \sum_{(\mathbf{T}, \theta) \in K} \frac{\varepsilon_{\mathbf{G}} \varepsilon_{\mathbf{T}} R_{\mathbf{T}}^{\mathbf{G}}(\theta)}{\langle R_{\mathbf{T}}^{\mathbf{G}}(\theta), R_{\mathbf{T}}^{\mathbf{G}}(\theta) \rangle}$$

for some unique geometric conjugacy class κ . We have $R_{\mathbf{T}}^G(\theta) \otimes \psi = R_{\mathbf{T}}^G(\theta(\psi|T))$ (see, e.g., [31, Proposition 12.6]), so

$$\rho \otimes \psi = \sum_{(\mathbf{T}, \theta) \in \kappa} \frac{\varepsilon_{\mathbf{G}} \varepsilon_{\mathbf{T}} R_{\mathbf{T}}^G(\theta(\psi|T))}{\langle R_{\mathbf{T}}^G(\theta), R_{\mathbf{T}}^G(\theta) \rangle}.$$

Since the distinct Deligne–Lusztig characters are orthogonal, the assumption $\rho \simeq \rho \otimes \psi$ implies that for any fixed $(\mathbf{T}, \theta) \in \kappa$, the pair $(\mathbf{T}, \theta(\psi|T))$ is also in the geometric conjugacy class κ .

Consider then the translation of this condition using the bijection between geometric conjugacy classes of pairs (\mathbf{T}, θ) and \mathbf{F}_q -rational conjugacy classes of semisimple elements in \mathbf{G}^* , the dual group of \mathbf{G} (see, e.g., [31, Proposition 13.12]; for instance, we have $\mathbf{G}^* = GL(n)$ if $\mathbf{G} = GL(n)$). Denote by s the conjugacy class corresponding to (\mathbf{T}, θ) .

The pair $(\mathbf{T}, \psi|T)$ corresponds to a *central* conjugacy class s' , because $\psi|T$ is the restriction of a *global* character of \mathbf{G} (see the proof of [31, Proposition 13.30]). Then, the definition of the correspondence shows that $(\mathbf{T}, \theta\psi|T)$ corresponds to the conjugacy class ss' , which is well-defined because s' is central. The assumption that (\mathbf{T}, θ) and $(\mathbf{T}, \theta\psi|T)$ are geometrically conjugate therefore means $ss' = s$, i.e., $s' = 1$, and clearly this means $\psi = 1$, as desired. \square

Remark 5.11 Here is a mnemonic device to remember the bounds for $A_{\infty}(G)$ in (5.1):⁵ among the representations of G , we have the principal series $R(\theta)$, parametrized by the characters of a maximal split torus, of which there are about q^r , and those share a common maximal dimension A . Hence

$$q^r A^2 \asymp \sum_{\theta} \dim(R(\theta))^2 \asymp |G| \sim q^d,$$

so A is of order $q^{(d-r)/2}$. In other words, we expect that in the formula $\sum \dim(\rho)^2 = |G|$, the principal series contributes a positive proportion.

The bound for $A_1(G)$ is also intuitive: there are roughly q^r conjugacy classes, and as many representations, and for a ‘positive proportion’ of them, the degree of the representation is of the maximal size given by $A_{\infty}(G)$.

⁵ Which explains why it seemed to the author to be a reasonable statement to look for . . .

6

Probabilistic sieves

The content of this chapter is a kind of warm-up to the next. Both involve applications of sieves where the siftable set is a general measure space (X, μ) , not simply a finite set with counting measure. The results described in this chapter may well be amenable to other proofs based on classical sieves, but this will not be the case in the next chapter. Moreover, alternative proofs may not be always possible if we go further along the route we describe . . .

The idea we want to pursue is to work with a given sieve setting (such as $\Psi = (\mathbf{Z}, \{\text{primes}\}, \mathbf{Z} \rightarrow \mathbf{Z}/\ell\mathbf{Z})$), using siftable sets which are probability spaces, given with a Y -valued random variable. Then we may look at the probability that the random variable lies in some sifted subset of Y , and as usual this may give information on the probability that the random variable satisfies certain properties which may be described or approached with sieve conditions. We pursue this in two ‘abelian’ cases here, before looking at non-abelian groups in the next chapter.

6.1 Probabilistic sieves with integers

Our first example is the analogue of the classical sieve of intervals of integers. Consider a probability space $(\Omega, \Sigma, \mathbf{P})$ (i.e., \mathbf{P} is a probability measure on Ω , which should not be confused with the sieving sets Ω_ℓ , with respect to a σ -algebra Σ ; see Appendix F for a survey of probabilistic language, for readers unfamiliar with it), and let $F = N : \Omega \rightarrow \mathbf{Z}$ be an integer-valued random variable. Then the triple (Ω, \mathbf{P}, N) is a siftable set, and given any sieving sets (Ω_ℓ) and prime sieve support \mathcal{L} , it is tautological that the probability of the associated sifted set in Ω is equal to

$$\mathbf{P}(N \in S(\mathbf{Z}, \Omega; \mathcal{L}^*)) = \mathbf{P}(\{\omega \in \Omega \mid \rho_\ell(N(\omega)) \notin \Omega_\ell, \text{ for all } \ell \in \mathcal{L}^*\}),$$

which is usually shortened to

$$\mathbf{P}(\rho_\ell(N) \notin \Omega_\ell \text{ for all } \ell \in \mathcal{L}^*)$$

(‘hiding’ the variable $\omega \in \Omega$, as usual in probability).

Note that this idea is one way of giving a precise meaning to natural quantities such as ‘the probability that an integer is squarefree’, if we are given natural integer-valued random variables. If the random variable is uniformly distributed among the integers $1 \leq n \leq N$, this becomes the usual density

$$\frac{|\{n \leq N \mid n \text{ is squarefree}\}|}{N}$$

and in general one is interested in the limit $N \rightarrow +\infty$ (in this case, the limit is well known to be $6/\pi^2$).

Exercise 6.1 Let N_λ be a random variable with a Poisson distribution of parameter λ , i.e., we have

$$\mathbf{P}(N_\lambda = k) = e^{-\lambda} \frac{\lambda^k}{k!}, \quad \text{for } k \geq 0.$$

Show that the probability that N_λ is squarefree (excluding 0) tends to $6/\pi^2$ as λ goes to ∞ .

We are more interested in random variables arising by means of a random walk on the integers. The philosophy is that such random walks provide the best approximation of the elusive ‘uniformly distributed random integer’ (since there is no translation-invariant probability on \mathbf{Z}). The continuous analogue is the idea that Brownian motion, in particular, gives the best understanding of what is a ‘random real number’. Although this idea seems natural enough, the only other work in this direction the author is aware of is a paper by Weber [131] concerning the behaviour of the number of divisors of a random integer obtained from a simple random walk.

A random walk (S_n) on \mathbf{Z} is simply a sequence of random variables (on some fixed probability space, as always) defined by

$$S_0 = 0, \quad \text{and} \quad S_{n+1} = S_n + X_{n+1} \tag{6.1}$$

where the increments (X_n) may also be arbitrary random variables. Of course, restrictions on (X_n) are usually imposed to conform with the intuition of a random walk. In particular, a common assumption is that (X_n) is a sequence of independent random variables, so that at each step the walker moves with no interference from the past (and the future!); or that the walk is a Markov

process, i.e., that S_{n+1} depends on S_0, \dots, S_n only through the value of S_n , or in conditional probability terms

$$\mathbf{P}(S_{n+1} = m \mid S_0 = m_0, \dots, S_n = m_n) = \mathbf{P}(S_{n+1} = m \mid S_n = m_n).$$

We will only consider the simplest case of the simple random walk on \mathbf{Z} , i.e., (S_n) is given by (6.1) and $(X_n)_{n \geq 1}$ is a sequence of independent random variables with centred Bernoulli distribution, namely

$$\mathbf{P}(X_n = \pm 1) = \frac{1}{2}.$$

These variables (S_n) give a natural sequence of siftable sets $(\Omega, \mathbf{P}, S_n)$. It turns out to be quite easy to estimate the corresponding large sieve constants, and the argument is a good illustration of the more sophisticated arguments to come in the next chapter.

Proposition 6.1 *Let (S_n) be a simple random walk on \mathbf{Z} . With notation as above, we have*

$$\Delta(S_n, \mathcal{L}) \leq 1 + \left| \cos\left(\frac{2\pi}{L^2}\right) \right|^n \sum_{m \in \mathcal{L}} m,$$

for $n \geq 1$ and for any sieve support \mathcal{L} consisting entirely of odd squarefree integers $m \leq L$. Here, since the dependency on the random variable component of the siftable set is the most important, we denote by $\Delta(S_n, \mathcal{L})$ instead of $\Delta(\Omega, \mathcal{L})$ the large sieve constant for the siftable set $(\Omega, \mathbf{P}, S_n)$.

Proof We will estimate the ‘exponential sums’, which in the current context, using probabilistic language, are the expectations

$$W(a, b) = \mathbf{E}\left(e\left(\frac{a_1 S_n}{m_1}\right) e\left(-\frac{a_2 S_n}{m_2}\right)\right) = \int_{\Omega} e\left(\frac{a_1 S_n}{m_1}\right) e\left(-\frac{a_2 S_n}{m_2}\right) d\mathbf{P},$$

for $m_1, m_2 \in \mathcal{L}$, $a_i \in (\mathbf{Z}/m_i \mathbf{Z})^\times$. Using the expression $S_n = X_1 + \dots + X_n$ for $n \geq 1$, independence, and the distribution of the X_i , we obtain straightforwardly

$$W(a, b) = \mathbf{E}\left(e\left(\frac{(a_1 m_2 - a_2 m_1) X_1}{m_1 m_2}\right)\right)^n = \left(\cos 2\pi \frac{a_1 m_2 - a_2 m_1}{m_1 m_2}\right)^n.$$

The conditions that m_i are odd, and that $(a_i, m_i) = 1$, imply that $|W(a, b)| = 1$ if and only if $a_1 = a_2$ and $m_1 = m_2$, and otherwise

$$|W(a, b)| \leq \left| \cos \frac{2\pi}{m_1 m_2} \right|^n.$$

Hence (see (2.8)), the large sieve constant is bounded by

$$\Delta(S_n, \mathcal{L}) \leq \max_{m_1, a_1} \left\{ 1 + \sum_{m_2} \sum_{a_2 \pmod{m_2}}^* \left| \cos \frac{2\pi}{m_1 m_2} \right|^n \right\} \leq 1 + \left| \cos \left(\frac{2\pi}{L^2} \right) \right|^n \sum_{m \in \mathcal{L}} m.$$

□

It is necessary to exclude even integers in this statement. The reason is simply that $S_n \pmod{2}$ is not equidistributed. Indeed, we clearly have $S_n \equiv n \pmod{2}$ for all n , so $\mathbf{P}(S_n \text{ is even}) = 0$ or 1 depending on whether n itself is even or odd. In probabilistic terms, the random walk is *periodic*.

Note that this difficulty is a consequence of the choice of the distributions of the increments (X_n) . Other distributions (taking values not restricted to ± 1) would avoid this. In particular, probably the simplest walk that avoids this problem is the one with independent increments X_n distributed according to

$$\mathbf{P}(X_n = \pm 1) = \frac{1}{4}, \quad \mathbf{P}(X_n = 0) = \frac{1}{2}, \quad (6.2)$$

or in other words, at each step the walker may decide to remain still with probability one-half, or to move in either of the two directions. Such walks are called *lazy* random walks.

Exercise 6.2 Prove analogues of the results below for the simple ‘lazy random walk’, without parity restrictions.

Corollary 6.2 *With notation as above, we have:*

(1) *For any sieving sets $\Omega_\ell \subset \mathbf{Z}/\ell\mathbf{Z}$ for ℓ odd, $\ell \leq L$, and $L \geq 3$, we have*

$$\mathbf{P}(S_n \in S(\mathbf{Z}, \Omega; L)) \leq \left(1 + L^2 \exp \left(-\frac{n\pi^2}{L^4} \right) \right) H^{-1}$$

where

$$H = \sum_{\substack{m \leq L \\ m \text{ odd}}}^b \prod_{\substack{\ell | m \\ \ell \text{ odd}}} \frac{|\Omega_\ell|}{\ell - |\Omega_\ell|}.$$

(2) *Let $\varepsilon > 0$ be given, $\varepsilon \leq 1/4$. For any odd $q \geq 1$, any a coprime with q , we have*

$$\mathbf{P}(S_n \text{ is prime and } \equiv a \pmod{q}) \ll \frac{1}{\varphi(q)} \frac{1}{\log n}$$

if $n \geq 2$, $q \leq n^{1/4-\varepsilon}$, the implied constant depending only on ε .

Note that (2) is Theorem 1.2 in the Introduction.

Proof For (1), we take \mathcal{L} to be the set of odd squarefree numbers $\leq L$ (so \mathcal{L}^* is the set of odd primes $\leq L$), and then since $\cos(x) \leq 1 - x^2/4$ for $0 \leq x \leq 2\pi/9$, the proposition gives

$$\Delta \leq 1 + L^2 \left(1 - \frac{\pi^2}{L^4}\right)^n \leq 1 + L^2 \exp\left(-\frac{n\pi^2}{L^4}\right),$$

and the result is a mere restatement of the large sieve inequality.

For (2), we have to change the sieve a little bit. Consider the sieve setting Ψ as above, *except* that for primes $\ell \mid q$, we take ρ_ℓ to be reduction modulo $\ell^{v(\ell)}$, where $v(\ell)$ is the ℓ -valuation of q . Take the siftable set (X, \mathbf{P}, S_n) , and the sieve support

$$\mathcal{L} = \{mm' \mid mm' \text{ squarefree, } (m, 2q) = 1, m \leq L/q \text{ and } m' \mid q\},$$

with \mathcal{L}^* still the set of odd primes $\leq L$.

Proceeding as in the proof of Proposition 6.1, the large sieve constant is bounded straightforwardly by

$$\Delta \leq 1 + \left| \cos \frac{2\pi}{L^2} \right|^n \sum_{\substack{m \leq L/q \\ (m, 2q)=1}} \sum_{m' \mid q} mq \leq 1 + \tau(q) q^{-1} L^2 \exp\left(-\frac{n\pi^2}{L^4}\right),$$

where $\tau(q)$ is the number of divisors of q , which satisfies $\tau(q) \ll q^\varepsilon$ for $q \geq 1$ and any $\varepsilon > 0$, the implied constant depending only on ε .

Finally, take

$$\Omega_\ell = \begin{cases} \{0\} & \text{if } \ell \nmid q, \\ \mathbf{Z}/\ell^{v(\ell)}\mathbf{Z} - \{a\} & \text{if } \ell \mid q. \end{cases}$$

If S_n is a prime number congruent to $a \pmod q$, then we have $S_n \in S(\mathbf{Z}, \Omega; \mathcal{L}^*)$, hence

$$\mathbf{P}(S_n \text{ is a prime} \equiv a \pmod q) \leq \mathbf{P}(S_n \in S(\mathbf{Z}, \Omega; \mathcal{L}^*)) \leq \Delta H^{-1}$$

where

$$H = \sum_{\substack{m \leq L/q \\ (m, 2q)=1}}^b \sum_{m' \mid q}^b \frac{\varphi^*(m')}{\varphi(m)}, \quad \text{with} \quad \varphi^*(n) = \prod_{\ell^v \parallel n} (\ell^v - 1),$$

where $\ell^v \parallel n$ means that ℓ^v divides n but ℓ^{v+1} does not. Now the desired estimate follows on taking $q \leq n^{1/4-\varepsilon}$ and $L = qn^\varepsilon$, using the classical lower bound (see, e.g. [11], [67, (6.82)])

$$\sum_{\substack{m \leq L/q \\ (m, 2q)=1}}^b \frac{1}{\varphi(m)} \geq \frac{\varphi(q)}{2q} \log L/q \gg \frac{\varphi(q)}{q} \log n$$

(the implied constant depending only on ε) together with the cute identity

$$\sum_{m' \mid q}^b \varphi^*(m') = q$$

which is trivially verified by multiplicativity. \square

Remark 6.3

- (1) It is important to keep in mind that, by the Central Limit Theorem, $|S_n|$ is usually of order of magnitude \sqrt{n} (see Appendix F), and precisely, S_n/\sqrt{n} converges in law to the normal distribution with variance 1 as $n \rightarrow +\infty$, so that for any real numbers $\alpha < \beta$, we have

$$\lim_{n \rightarrow +\infty} \mathbf{P}\left(\alpha \leq \frac{S_n}{\sqrt{n}} \leq \beta\right) = \frac{1}{\sqrt{2\pi}} \int_{\alpha}^{\beta} e^{-t^2/2} dt.$$

So the estimate $\Delta \leq 1 + L^2 \exp(-n\pi^2/L^4)$, which gives a non-trivial result in applications as long as, roughly speaking, $L \leq n^{1/4}/(\log n)^{1/4}$, compares well with the classical large sieve for integers $n \leq N$, where $\Delta \leq N - 1 + L^2$, which is non-trivial for $L \leq \sqrt{N}$.

- (2) The second part is an analogue of the Brun–Titchmarsh inequality, namely (in its original form)

$$\pi(x; q, a) \ll \frac{1}{\varphi(q)} \frac{x}{\log x} \quad (6.3)$$

for $x \geq 2$, $(a, q) = 1$ and $q \leq x^{1-\varepsilon}$, the implied constant depending only on $\varepsilon > 0$. However, from the previous remark we see that it is weaker than could be expected, namely $q \leq n^{1/4-\varepsilon}$ would have to be replaced by $q \leq n^{1/2-\varepsilon}$. Here we have exploited the flexibility of the sieve setting and sieve support. For a different use of this flexibility, see Chapter 8.¹

It would be quite interesting to know if the extension to $q \leq n^{1/2-\varepsilon}$ holds. The point is that if we try to adapt the classical method, which is to sieve for those k , $1 \leq k \leq x/q$, such that $qk + a$ is prime, we are led to some interesting and non-obvious (for the author) probabilistic issues; indeed, if $S_n \equiv a \pmod{q}$, the (random) integer k such that $S_n = kq + a$ can be described as follows (using some standard properties of the simple random walk on \mathbf{Z} , e.g., the determination of the probability that such a walk first

¹ We want to point out here that the possibility of using a careful non-obvious choice of \mathcal{L} in the large sieve was exploited by D. Zywinia in his preprint ‘The large sieve and Galois representations’. Making such a choice is also, to a large degree, the very point of combinatorial sieves, starting with V. Brun’s work, though the emphasis there is very different (see Appendix A for a few words and references about this).

reaches $-b$ or a for given integers $a, b \geq 1$): we have $k = T_N$ where N is a random variable

$$N = |\{m \leq n \mid S_m \equiv a \pmod{q}\}|$$

and (T_i) is a random walk with initial distribution given by

$$\mathbf{P}(T_0 = 0) = 1 - \frac{a}{q}, \quad \mathbf{P}(T_0 = -1) = \frac{a}{q},$$

and independent identically distributed increments $V_i = T_i - T_{i-1}$ such that

$$\mathbf{P}(V_i = 0) = 1 - \frac{1}{q}, \quad \mathbf{P}(V_i = \pm 1) = \frac{1}{2q}.$$

So what is needed is to perform sieve on the siftable set

$$(\{S_n \equiv a \pmod{q}\}, \mathbf{P}, T_N)$$

where the length N of the auxiliary walk is random. Note, at least, that if we look at the same problem with $(\Omega, \mathbf{P}, T_i)$ for a fixed i , then we easily get by sieving that

$$\mathbf{P}(qT_i + a \text{ is prime}) \ll \frac{1}{\varphi(q)} \frac{1}{\log i}$$

for all $q \leq i^{1/2-\varepsilon}$, $\varepsilon > 0$, the implied constant depending only on ε .

- (3) Obviously, it would be very interesting to derive lower bounds or asymptotic formulas for $\mathbf{P}(S_n \text{ is prime})$ for instance, and for other analogues of classical problems of analytic number theory. Note that it is tempting to attack the problems with ‘local’ versions of the Central Limit Theorem and summation by parts to reduce to the purely arithmetic deterministic case. This is unlikely to be possible in other cases however, such as in the next section.

Before concluding with an exercise that follows the trail of our *fil rouge*, we remark that this probabilistic point of view should not be mistaken with ‘probabilistic models’ of integers (or primes), such as Cramer’s model: the values of the random variables we have discussed are perfectly genuine integers.²

Exercise 6.3 Consider the following random walk (P_k) on the set of monic polynomials of degree $d \geq 1$ with integral coefficients: at each step, a degree i , $0 \leq i \leq d-1$, is chosen uniformly randomly, and the coefficient of X^i is either increased by 1 or decreased by 1, according to a centred symmetric Bernoulli distribution. (Of course, each step is independent from all others.)

² To give a caricatural example, if it were possible to show that, for some sequence of random variables N_n distributed on disjoint subsets of integers, the probability $\mathbf{P}(N_n \text{ and } N_n + 2 \text{ are both primes})$ is always strictly positive, then the twin-prime conjecture would follow.

(1) Show that

$$\mathbf{P}(P_k \text{ is reducible}) \ll k^{-1/4} (\log k)^2$$

for $k \geq 2$, the implied constant depending only on d , by following Gallagher's approach (see Theorem 4.2) and setting up a d -dimensional probabilistic large sieve.

(2) Show that the set X_d of reducible polynomials (monic of degree d) is *recurrent* for this walk, i.e., almost surely there are infinitely many k such that $P_k \in X_d$. [**Hint:** Show that the random variable $P_k(1)$ itself follows a simple random walk on \mathbf{Z} , and hence, almost surely, $P_k(1)$ is zero infinitely often (see Appendix F).] This last part should be contrasted with Theorem 1.3. (Note also that for $d \geq 3$, the random walk (X_k) itself is transient, in the sense that the probability of returning infinitely often to the same polynomial is zero.)

6.2 Some properties of random finitely presented groups

The second example in this chapter will use a probabilistic sieve in integer matrices to study some asymptotic properties of certain types of random finitely presented groups.

There exist different notions of 'random groups' in the literature; the one we consider here is the one described by Dunfield and Thurston in [34, Section 3], which they use as a basis for comparison with fundamental groups of certain types of random 3-manifolds. We will discuss some applications of the sieve to those particular groups in the next chapter, which serves as 'forward' motivation for this section (see Proposition 7.19 and the surrounding discussion in Section 7.6, which the reader may wish to look at quickly after finishing reading this chapter).

Let $g \geq 2$ be an integer, and let F_g be the free group on g generators a_1, \dots, a_g . We consider groups G obtained as quotients of F_g by normal subgroups generated by g words³ of length k in the generators a_i and their inverses a_i^{-1} . More precisely, we consider all $(2g)^{gk}$ presentations

$$G_k = \langle F_g \mid w_{1,k}, \dots, w_{g,k} \rangle$$

³ We consider 'balanced' presentations, i.e., with as many relations as there are generators because, as explained in [34], this is the critical case for the size of the abelianization.

where each $w_{i,k}$ is such a word. In probabilistic terms, each $W_k = (w_{1,k}, \dots, w_{g,k})$ is the k -th step of a sequence of g independent random walks on F_g with uniformly chosen independent steps in $S = \{a_i, a_i^{-1}\}$.

As in [34, 3.14],⁴ and as in Proposition 7.19, we consider the abelianization $G_k/[G_k, G_k]$ of the random group G_k , and we show that on the one hand, it is a *finite* group with high probability, but that its order is also large with high probability, because it has non-zero p -primary parts for many primes p .

Proposition 6.4 *Let $g \geq 2$ be an integer, and for $k \geq 1$, let G_k be the above random group.*

(1) *We have*

$$k^{-g/2} \ll \mathbf{P}(G_k/[G_k, G_k] \text{ is finite}) \ll k^{-1/2} \log k,$$

for $k \geq 1$, where the implied constants depend only on g . If $g = 2$, then we have more precisely

$$\mathbf{P}(G_k/[G_k, G_k] \text{ is finite}) \asymp k^{-1}.$$

(2) *We have*

$$\mathbf{P}\left(|G_k/[G_k, G_k]| < (\log k)^{\beta \log \log k}\right) \ll \frac{1}{\log \log k}$$

for $k \geq 3$, where $\beta > 0$ and the implied constant depends only on g .

Proof By definition of G_k , the abelian group $G_k/[G_k, G_k]$ is the quotient of the free abelian group $\mathbf{Z}^g = F_g/[F_g, F_g]$ with basis (e_1, \dots, e_g) by the subgroup generated by the ‘abelianized’ relations $v_{i,k}$ which are the image of $w_{i,k}$ in \mathbf{Z}^g : if

$$w = \prod_{1 \leq j \leq k} a_{i_j}^{\delta_j}$$

(with $\delta_j \in \{-1, 1\}$) is a word of length k , we have

$$v = \sum_{1 \leq i \leq g} \left(\sum_{i_j = i} \delta_j \right) e_i \in \mathbf{Z}^g.$$

Let M_k be the $g \times g$ integral matrix with columns given by the column vectors $v_{i,k}$. Then the theory of abelian groups of finite type states that $G_k/[G_k, G_k]$ is finite if and only if $\det M_k \neq 0$, and then $|G_k/[G_k, G_k]| = |\det M_k|$. Now the point is that the sequence of matrix-valued random variables (M_k) is obtained

⁴ Dunfield and Thurston allow the steps to be the identity, to avoid periodicity issues. This case can also be treated, with the same results, but it introduces complications in the notation so we select this simpler random walk.

by a random walk on the additive group of integral $g \times g$ matrices, which is isomorphic to \mathbf{Z}^{g^2} , with the following description: $M_{k+1} = M_k + A_{k+1}$ where (A_k) is a sequence of independent matrix-valued random variables, uniformly supported on the set T of matrices with exactly one non-zero entry per column, which is equal to either 1 or -1 ; note that $|T| = (2g)^g$ (T does not generate the group of integral matrices, since each element of T has the property that the sum of entries in each column is constant modulo 2, but this is not a problem). We are thus led to study the distribution of the determinant of the random integral matrices (M_k) . Note of course that all of this first part is contained in [34].

To bound from above the probability that G_k has infinite abelianization, we may simply use equidistribution in finite quotients, since a non-invertible matrix reduces to a non-invertible matrix modulo any prime: for any odd prime ℓ , we find that

$$\begin{aligned} \mathbf{P}(\det M_k = 0) &\leq \mathbf{P}(M_k \pmod{\ell} \notin GL(g, \mathbf{F}_\ell)) \\ &= 1 - \frac{|GL(g, \mathbf{F}_\ell)|}{\ell^{g^2}} + O\left(\ell^{(g^2-1)/2} \exp\left(-\frac{gk\pi^2}{\ell^2}\right)\right) \end{aligned}$$

with an absolute implied constant, using Remark 2.14 and the estimates for exponential sums in the proof of Lemma 6.6 below.

Since

$$1 - \frac{|GL(g, \mathbf{F}_\ell)|}{\ell^{g^2}} = 1 - \prod_{1 \leq i \leq g} (1 - \ell^{-i}) = \frac{1}{\ell} + O\left(\frac{1}{\ell^2}\right) \quad (6.4)$$

for $\ell \geq 2$, the implied constant depending only on g , we have

$$\mathbf{P}(\det M_k = 0) \leq \frac{1}{\ell} + O\left(\frac{1}{\ell^2}\right) + O\left(\ell^{(g^2-1)/2} \exp\left(-\frac{gk\pi^2}{\ell^2}\right)\right).$$

If we take ℓ such that $\ell \asymp \sqrt{g^{-1}k(\log k)^{-1}}$, we obtain

$$\mathbf{P}(\det M_k = 0) \ll k^{-1/2} \log k$$

with an implied constant depending only on g ; as usual, if k is too small for such a prime ℓ to exist, the estimate is trivial when the implied constant is large enough.

To get a lower bound, on the other hand, we use the following trivial observation: we have $\det M_k = 0$ if the first two columns of M_k are identical. Now each of those two columns is obtained by a simple random walk on \mathbf{Z}^g , and the two walks are independent. It is then a fairly simple fact that the probability that the two random walks coincide at the k -th step is of size $\gg k^{-g/2}$ for $k \geq 1$ (see (F.3) in Appendix F), the implied constant depending on g , as stated.

In the case $g = 2$, we can refine the bound by foregoing any reduction: the probability that the integral matrix M_k be singular is simply the probability that its two rows are proportional, i.e., it is

$$\mathbf{P}(nX_k = mY_k \text{ for some } (n, m) \neq (0, 0) \in \mathbf{Z}^2)$$

for two independent simple random walks on \mathbf{Z}^2 (the column vectors of the matrix). This is of size k^{-1} as explained in Section F.4 (this result was shown by D. Khoshnevisan). Moreover, the set of singular matrices is recurrent in this case: indeed, already the set of matrices with two identical columns is recurrent (see, again, Appendix F).

We now come to the second part of the proposition, which is really where sieve is used. Note first that if $G_k/[G_k, G_k]$ is finite, then it has order divisible by a prime ℓ if and only if $\det M_k = 0 \pmod{\ell}$. To detect such occurrences, we apply the dual sieve (2.13), with Ω_ℓ the complement of $GL(g, \mathbf{F}_\ell)$ in the group of $g \times g$ matrices, for all odd primes ℓ . This means that we consider the probabilistic sieve setting on \mathbf{Z}^{g^2} with the reduction maps modulo primes, and the siftable set associated to the random variable M_k .

By Lemma 6.6 below, the large sieve constant satisfies

$$\Delta(M_k, \mathcal{L}) \leq 1 + L^{g^2+1} \exp\left(-\frac{gk\pi^2}{L^4}\right)$$

for the prime sieve support and sieve support both consisting of odd primes $\ell \leq L$, so that the dual sieve leads to the inequality

$$\mathbf{E}\left(\left(P(M_k, L) - P(L)\right)^2\right) \leq \left(1 + L^{g^2+1} \exp\left(-\frac{gk\pi^2}{L^4}\right)\right) P(L)$$

where $P(M_k, L)$ is the number of odd primes $\ell \leq L$ such that $\det M_k \equiv 0 \pmod{\ell}$, and

$$P(L) = \sum_{3 \leq \ell \leq L} \frac{|\Omega_\ell|}{\ell^{g^2}} = \sum_{3 \leq \ell \leq L} \frac{1}{\ell} + O(1) \gg \log \log L$$

for $L \geq 3$, the implied constant depending on g , by (6.4).

Let $L = (k/(g \log k))^{1/4}$ (if this is ≥ 3 , otherwise, we can increase the implied constant at the end as usual); it follows by positivity that for some constant $c > 0$ (depending on g), we have

$$\mathbf{P}(P(M_k, L) < c \log \log k) \ll \frac{1}{\log \log k}.$$

This conclusion is actually more precise than the next step, and indeed it is best possible, because it is clear that the determinant of M_k is at most of polynomial size in k (each coefficient of M_k is at most k , so the absolute value

of the determinant is trivially $\leq g!k^g$, and (if non-zero) it can only have a bounded number of additional prime divisors $\ell > L$ for L as above (the bound depending on g , of course). In particular, note that this shows that the expected value of the size of the torsion subgroup of $G_k/[G_k, G_k]$ is $\leq g!k^g$ for $k \geq 1$, i.e., it grows at most polynomially.

To obtain the inequality in (2), we simply observe that if $P(M_k, L) \geq c \log \log k$, we have

$$|G_k/[G_k, G_k]| \geq \prod_{3 \leq \ell \leq Q} \ell$$

where Q is the $[c \log \log k]$ -th prime. The Chebychev estimates prove that this is at least $(\log k)^{\beta \log \log \log k}$, for some $\beta > 0$ depending on g (because c does); see the end of the proof of Proposition 7.19 for details, if needed. \square

Remark 6.5 It seems likely that in fact the size of $G_k/[G_k, G_k]$ is at least a power of k with probability tending to 1 as $k \rightarrow +\infty$; showing this requires knowing that the primes dividing this number are not almost always ‘too’ small.

As in the case of the elliptic sieve, one may even speculate whether the determinant of M_k behaves as a ‘typical’ integer in other manners, for instance (optimistically), does there exist an analogue of the Erdős–Kac theorem?

Another question, related to the first part, is to know which of the upper and lower bounds is closer to the truth. In particular, although both together show polynomial decay of the probability that G_k has infinite abelianization, the discrepancy does not allow to say whether the set of groups with infinite abelianization is transient or not (the upper bound is too large to apply the Borel–Cantelli lemma, as we will do in the next chapter for similar problems). However, it seems likely that the lower bound is closer to the truth, and therefore that the answer is that this set is transient, except for $g = 1$ or $g = 2$ (where we know that the set is indeed recurrent).

For further remarks, see the discussion after Proposition 7.19.

Finally, here is the computation of the large sieve constant used in the proof:

Lemma 6.6 *With notation as in the proof of Proposition 6.4, the large sieve constant for the siftable set M_k and the sieve support of odd primes $\ell \leq L$ satisfies*

$$\Delta(M_k, \mathcal{L}) \leq 1 + L^{g^2+1} \exp\left(-\frac{gk\pi^2}{L^4}\right)$$

for $k \geq 1$.

Proof This is close to what was done in Section 6.6 (which is more or less the case $g = 1$). We use the additive characters of $\mathbf{F}_\ell^{g^2}$ as basis elements to estimate

the exponential sums; in terms of matrices, these characters are given by

$$m \mapsto e\left(\frac{\text{Tr } mn}{\ell}\right)$$

where n runs over $g \times g$ matrices with coefficients in \mathbf{F}_ℓ .

So, fix two odd primes $\ell, \ell' \leq L$, two non-zero matrices n, n' (of size g) with coefficients in \mathbf{F}_ℓ and $\mathbf{F}_{\ell'}$ respectively, and let

$$W(n, n') = \mathbf{E}\left(e\left(\frac{\text{Tr}(M_k n)}{\ell} - \frac{\text{Tr}(M_k n')}{\ell}\right)\right).$$

If $(\ell, n) = (\ell', n')$, we have $W(n, n') = 1$, obviously, so suppose this is not the case. Since $M_k = A_1 + \dots + A_k$ and the steps are independent, we obtain

$$W(n, n') = \mathbf{E}\left(e\left(\frac{\text{Tr}(An)}{\ell} - \frac{\text{Tr}(An')}{\ell}\right)\right)^k$$

just as in the proof of Proposition 6.1, where A is distributed as the steps of the random walk, i.e., it is a random variable with values in the set T described at the beginning of the proof of Proposition 6.4, each $t \in T$ having identical probability $(2g)^{-g}$.

To simplify notation, we look at

$$\mathbf{E}(e(\text{Tr } vA)) = \frac{1}{|T|} \sum_{t \in T} e(\text{Tr } vt)$$

where $v = (v_{i,j})$ is a real-valued matrix ($v = (\ell'n - \ell n')/\ell\ell'$ in the application). We rewrite the sum over $t \in T$ by putting outside the average over the g^g possible choices of one position (i_1, \dots, i_g) in each column (each i_j is the index of the row where the non-zero entry is found in the j -th column), and inside the average over the 2^g choices of elements $(\pm 1, \dots, \pm 1)$ to be placed in those positions. The inner sum becomes

$$\frac{1}{2^g} \prod_{1 \leq k \leq g} (e(v_{i_k, k}) - e(-v_{i_k, k})) = \prod_{1 \leq k \leq g} \cos 2\pi v_{i_k, k}.$$

With the assumptions on ℓ and ℓ' , n and n' , for the given v we have

$$\left| \prod_{1 \leq k \leq g} \cos 2\pi v_{i_k, k} \right| \leq \left| \cos \frac{2\pi}{L^2} \right|^g,$$

and hence, as in the proof of Corollary 6.2, we get

$$|W(n, n')| \leq \left| \cos \frac{2\pi}{L^2} \right|^{gk} \leq \exp\left(-\frac{gk\pi^2}{L^4}\right),$$

and by (2.8) we have

$$\Delta \leq 1 + \sum_{\ell \leq L} \ell^{g^2} \exp\left(-\frac{gk\pi^2}{L^4}\right) \leq 1 + L^{g^2+1} \exp\left(-\frac{gk\pi^2}{L^4}\right).$$

□

7

Sieving in discrete groups

7.1 Introduction

This chapter, which may be the most innovative in this book, reflects the outcome of a number of different important mathematical ideas. Most of them are related to number theory, but as we will see, both the tools involved in making the sieve apply and its potential applications go far beyond.

The basic motivation is that any discrete set with interesting structure can be investigated by ideas that are related to sieve. The object we consider here, for the most part, is a discrete finitely generated group G (see the last remark in this introductory section for some words on another variant arising from the ongoing work of Bourgain, Gamburd and Sarnak). Of course, the simplest such group is undoubtedly \mathbf{Z} , which recovers the classical sieve setting. If we stick to groups with an arithmetic flavour, it seems natural, however, to consider for instance the modular group $SL(2, \mathbf{Z})$, which intervenes prominently in both analytic and algebraic number theory, and then more generally $SL(n, \mathbf{Z})$, $n \geq 2$, or $Sp(2g, \mathbf{Z})$, $g \geq 1$. For each of these groups, there is an obvious reduction map ρ_ℓ modulo a prime ℓ , with image a finite group, which is indeed a finite group of Lie type (such as were considered in Chapter 5), and this gives a sieve setting $(G, \{\text{primes}\}, (\rho_\ell))$. In fact, any ‘arithmetic group’ is a natural target for sieving but, for simplicity, we will keep to the most concrete cases.

The first definite problem to keep in mind, at least from the point of view of analytic number theory, is probably the following: is it true that given a matrix $g \in SL(n, \mathbf{Z})$ or $Sp(2g, \mathbf{Z})$ with ‘norm’ bounded by some quantity T , the ‘probability’ that the characteristic polynomial $\det(T - g) \in \mathbf{Z}[T]$ of g be reducible tends to 0 as $T \rightarrow +\infty$? This is a very natural question, considering that the corresponding fact holds for polynomials of given degree with bounded coefficients, and indeed, looking at the result of Gallagher (see Theorem 4.2), it is understandable to wish for the sharpest possible result in

this direction. However (although we started considering this independently), as far as we know, the first public mention of this question is to be found in Rivin's paper [108, Conjecture 8].

Exercise 7.1 Looking at unimodular (or symplectic) matrices is what makes the problem difficult; indeed, let $n \geq 2$ be an integer, $N \geq 1$, and define

$$R_n(N) = |\{g = (g_{i,j}) \in M(n, \mathbf{Z}) \mid |g_{i,j}| \leq N, \text{ and } \det(T - g) \text{ is reducible}\}|,$$

$$R'_n(N) = |\{g \in GL(n, \mathbf{R}) \cap M(n, \mathbf{Z}) \mid \|g\| \leq N, \text{ and } \det(T - g) \text{ is reducible}\}|$$

(where $M(n, \mathbf{Z})$ is the ring of $n \times n$ matrices with integral coefficients, with no condition on the determinant, and $\|g\|$ is defined below in (7.1)). Show that

$$|R'_n(N)| \leq |R_n(N)| \ll N^{n^2-1/2}(\log N)$$

for $N \geq 2$, where the implied constant depends only on n . [Hint: Use the n^2 -dimensional classical large sieve and argue as in Gallagher's Theorem 4.2.]

Now, on more careful consideration of this type of idea, it quickly appears that one may in fact consider different types of siftable sets, and obtain problems with distinctly different (more analytic, combinatorial or probabilistic) flavour, depending on which is chosen:

- The most analytic type of siftable sets are the finite sets such as

$$X = \{g \in SL(n, \mathbf{Z}) \mid \|g\| \leq T\}$$

with the counting measure, and identity mapping $X \rightarrow G$. Here the norm $\|g\|$ of a matrix might be any fixed norm; a natural one to consider is

$$\|g\| = \left(\sum_{1 \leq i,j \leq n} |g_{i,j}|^2 \right)^{1/2}, \quad \text{for } g = (g_{i,j}) \in GL(n, \mathbf{R}), \quad (7.1)$$

which has the property that $\|gh\| = \|hg\| = \|g\|$ for any orthogonal matrix $h \in O(n, \mathbf{R})$.

Here the equidistribution approach leads to hyperbolic lattice point problems (in the case $n = 2$), and generalizations of those for $n \geq 3$. The issue of uniformity with respect to q when looking at the principal congruence subgroups $\Gamma(q)$ is the main issue, compared with the results which are available in the literature, and which are quite complete – and work in much greater generality – for an individual subgroup. For instance, the original work of Duke, Rudnick and Sarnak [33] gives individual equidistribution in $SL(n, \mathbf{Z})$ modulo a prime, using methods of harmonic analysis. There have been generalizations and alternative treatments using ergodic-theoretic

methods, for instance by Eskin, Mozes, McMullen, Shah and others (see, e.g. [37]). However, a uniform treatment, as required for an efficient application of the large sieve, is not so obvious. This is however likely to follow soon from ongoing work of Sarnak and Nevo, building on the methods of Duke, Rudnick and Sarnak (it would be very interesting to have similar uniformity from ergodic methods, but that seems to be a very difficult question).

The two (or two and a half, as will be seen) other settings are in fact suitable for much more general groups (in principle). Indeed, let G be an arbitrary finitely generated discrete group and suppose a finite generating set S of G is given (which we assume to be symmetric, i.e., such that $S^{-1} = S$); note that it is well known that $SL(n, \mathbf{Z})$, $n \geq 1$, and $Sp(2g, \mathbf{Z})$, $g \geq 1$, are finitely generated, as will be recalled below. Then we may be interested in the following types of siftable sets associated to G (note that we have not completely defined a sieve setting in full generality, but the point is partly that the siftable set may suggest itself *before* – or independently of – the ‘reduction maps’ ρ_ℓ):

- The set X of elements $g \in G$ with word-length metric $\ell_S(g)$ at most N , for some integer $N \geq 1$, i.e., the set of those elements $g \in G$ that can be written as

$$g = s_1 \cdots s_k$$

with $k \leq N$, $s_i \in S$ for $1 \leq i \leq k$. To make a siftable set, we would take here counting measure with identity map $X \rightarrow G$.

- The set W of words $s_1 s_2 \cdots s_N$ of length N in the alphabet S , for some integer $N \geq 1$, with the map $w \mapsto F_w$ being this time the obvious ‘evaluation’ in G of the word w . Note that of course F may not be injective; also $|W| = |S|^N$, which is thus exponentially growing as a function of N .
- More generally, the siftable set $(W, \text{counting measure}, F_w)$ can be interpreted as a specific generalization of the probabilistic context of the previous chapter: the values of the words in W may be seen as the result, after N steps, of a random walk on G obtained by starting from the origin and ‘walking’ by multiplying on the right at each step by a uniformly chosen, randomly selected, element of S (independently of any other step). Now consider instead a fairly general random walk of this type, namely let $(\Omega, \Sigma, \mathbf{P})$ be a probability space, and suppose that a sequence (ξ_k) , $k \geq 1$, of independent S -valued random variables is given. Then define the corresponding left-invariant random walk (X_k) on G by

$$X_0 = 1 \in G, \quad X_{k+1} = X_k \xi_{k+1} \text{ for all } k \geq 0.$$

This yields a natural sequence $(\Omega, \mathbf{P}, X_k)$ of probabilistic siftable sets. The simplest case, as with the simple random walk on \mathbf{Z} , is to assume that ξ_k is uniformly distributed:

$$\mathbf{P}(\xi_k = s) = \frac{1}{|S|}, \text{ for all } s \in S,$$

and in that case, as we observed, this siftable set is equivalent with the set W of words of length k : we have

$$\mathbf{P}(X_k = g) = \frac{|\{w \in W \mid F_w = g\}|}{|W|}$$

for any $g \in G$, hence any result may be stated for one or the other formulation. Although the set of words is more concrete in a sense, avoiding probabilistic language, some results are definitely better phrased probabilistically (see Theorem 1.3).

Also, we will wish to vary the distribution of the factors ξ_k of the random walk because in some applications an assumption of uniform distribution is not necessarily valid (see the proof of Proposition 7.17). For instance, another natural type of random walk is given, when S does not contain 1, by steps ξ_k where

$$\mathbf{P}(\xi_k = s) = \frac{1}{2|S|} \quad \text{for } s \in S, \quad \mathbf{P}(\xi_k = 1) = \frac{1}{2}, \quad (7.2)$$

(this generalizes (6.2)). This type of ‘lazy’ walk will avoid periodicity problems similar to the problem with the simple random walk on \mathbf{Z} mentioned after Proposition 6.1.

After listing those types of siftable sets, many readers will probably feel that either the analytic siftable sets (and their attending hyperbolic lattice-point problems), or the balls in word-length metric, are the most natural and interesting. However, we will consider below the probabilistic sieves. The reason is that those are in fact the easiest to deal with, and that they can be handled very transparently by invoking some important and well-established aspects of harmonic analysis on groups, namely Property (T) of Kazhdan or Property (τ) of Lubotzky. It seems clear that dealing with the other sieve settings will necessarily involve the same ingredients, and others, but there will be additional complications and it is by no means clear that the same generality can be achieved. Another good reason to study random walks is that this can provide rigorous results to develop an intuition of what a ‘typical’ element looks like, which is very useful in situations where this is not (yet) clear (see the discussion in [34] concerning the case of random 3-manifolds, which we will discuss briefly in Section 7.6).

Before embarking on this study, let us mention two references that the reader may find useful (as we did): P. de la Harpe gives a highly readable and entertaining account of many topics of ‘geometric’ group theory in [57] and L. Saloff-Coste has a very clear and enlightening survey of random walks on (mostly finite) groups in [111]. We also mention that, although there is an extensive theory of random walks on groups in general (see, e.g. [132], and the pioneering work of Furstenberg [45]), we will not use any of this. It would be interesting to find deeper interactions between this theory and the sieve techniques.

Remark 7.1 There are certainly other types of sieve settings that may be interesting. The work of Bourgain, Gamburd and Sarnak (see [14, 15] and Sarnak’s slides for the Rademacher Lectures [113]) is based on the following: consider a finitely generated group Γ which is a discrete subgroup of a matrix group over \mathbf{Z} , acting on an affine algebraic variety V/\mathbf{Z} . Then the sieve setting is $(\Gamma \cdot v, \{\text{primes}\}, \rho_\ell)$ where $\Gamma \cdot v$ is the orbit of a fixed element $v \in V(\mathbf{Z})$, and ρ_ℓ is the reduction map to the finite orbit of the reduction in $V(\mathbf{F}_\ell)$ (with uniform density). The siftable set is a subset Y of the orbit defined by the images of elements of Γ of bounded word-length or bounded norm, with counting measure and identity map.

7.2 Random walks in discrete groups with Property (τ)

We now consider the third (hence also the second) type of probabilistic siftable set $(\Omega, \mathbf{P}, X_k)$ for a finitely generated discrete group G , with a finite symmetric generating set S .

As mentioned in the previous section, we have not identified a specific sieve setting to go with our group, and indeed for the moment we will simply assume that we are given some family (ρ_ℓ) , $\ell \in \Lambda$, of surjective homomorphisms

$$\rho_\ell : G \rightarrow G_\ell$$

onto finite groups. We assume moreover for simplicity that the steps (ξ_k) of the random walk (X_k) are symmetric and identically distributed so that

$$\mathbf{P}(\xi_k = s) = \mathbf{P}(\xi_k = s^{-1}) = p(s) \in [0, 1], \text{ for all } s \in S.$$

Obviously, the random variable X_k lies (almost surely) in the group generated by those $s \in S$ for which $p(s) > 0$; if this group is strictly smaller than G , we might as well have started from it (and the generating set obtained by removing those s with $p(s) = 0$), so we also assume that $p(s) > 0$ for all s .

We will now derive a bound for the large sieve constant for the group sieve setting $(G, \Lambda, (\rho_\ell))$, or the associated conjugacy sieve, under some analytic conditions on the group G and the family (ρ_ℓ) . We could use the simple equidistribution approach (see Section 3.4). However, it is really cleaner and more efficient to estimate the exponential sums $W(\pi, \tau)$ or $W(\varphi_{\pi, e, f}, \varphi_{\tau, e', f'})$ of (3.2), and then apply (2.8). Also, we consider a case where the starting point X_0 of the random walk is not necessarily the identity $1 \in G$, but may be any G -valued random variable, supported for simplicity on finitely many elements, and independent of the steps (ξ_k) . (We will have the opportunity to use this more general case in Proposition 7.11.)

The crucial ingredient is the so-called Property (τ) , and in fact the argument is almost tautological given this assumption. We will recall the definition in the course of the proof, and we refer to Appendix D for some more details, or, e.g., [91, Section 4.3] or [93] for more complete surveys.

Proposition 7.2 *Let G be a finitely generated discrete group, I an arbitrary index set and (N_i) a family of finite index normal subgroups of G for $i \in I$.*

Let $S = S^{-1}$ be a symmetric finite generating set of G , and let (X_k) , $k \geq 0$, denote a left-invariant symmetric random walk on G given by an initial step X_0 supported on a finite set $T \subset G$, and by $X_{k+1} = X_k \xi_{k+1}$ with independent steps (ξ_k) identically distributed with

$$\mathbf{P}(\xi_k = s) = \mathbf{P}(\xi_k = s^{-1}) = p(s) > 0, \quad \text{for all } s \in S,$$

and moreover chosen so that X_0 is independent of ξ_k .

Assume that:

- *The group G has Property (τ) with respect to the family (N_i) of finite index subgroups.*
- *There is a word $r = s_1 \cdots s_c$ in the alphabet S of odd length c , such that $s_1 \cdots s_c = 1$ in G ; for instance, this holds if $1 \in S$ or if there is no non-trivial homomorphism $G \rightarrow \mathbf{Z}/2\mathbf{Z}$.*

Then there exists $\eta > 0$ such that for any finite-dimensional representation $\pi : G \rightarrow GL(V)$ with $\text{Ker } \pi \supset N_i$ for some i , either there exists a non-zero $v \in V$ invariant under G , or we have

$$\left| \mathbf{E}(\langle \pi(X_k)e, f \rangle) \right| \leq \|e\| \|f\| \exp(-\eta k) \quad (7.3)$$

for any vectors e, f in the space of π and any $k \geq 0$, where $\langle \cdot, \cdot \rangle$ is a G -invariant inner product on V .

In particular, either there is a non-zero invariant vector, or we have

$$\left| \mathbf{E}(\text{Tr } \pi(X_k)) \right| \leq (\dim \pi) \exp(-\eta k), \quad (7.4)$$

for any $k \geq 0$.

The constant η depends only on the distribution of the steps of the walk, the (τ) -constant for (G, S, N_i) and the length c of the relation r ; indeed one can take

$$\eta = \min \left(\log \frac{1}{1 - \frac{\kappa p^+}{2}}, \log \frac{1}{1 - \frac{2p^+}{c^2}} \right) \geq p^+ \min \left(\frac{\kappa}{2}, \frac{2}{c^2} \right), \quad (7.5)$$

where $\kappa = \kappa(G, S, (N_i))$ is the (τ) -constant for G , and $p^+ = \min p(s)$ is the smallest probability of a generating element.

In particular, for any integer $N \geq 1$, let $W = W_N$ denote the set of words of length N in the alphabet S , and let F_w denote the value of the word w in G . Under the assumptions above, if V contains no non-zero invariant vectors, we have

$$\left| \sum_{w \in W} \langle \pi(F_w)e, f \rangle \right| \leq \|e\| \|f\| |W|^{1-\alpha}, \quad (7.6)$$

$$\left| \sum_{w \in W} \text{Tr } \pi(F_w) \right| \leq (\dim \pi) |W|^{1-\alpha} \quad (7.7)$$

with $\alpha = \eta / \log |S|$, η being computed with $p^+ = 1/|S|$.

This proposition should also be compared with [111, Theorem 6.15]. In a sieve setting, it will be applied with $I = S(\Lambda)$ and $N_m = \text{Ker}(\rho_m)$ for $m \in S(\Lambda)$, with π replaced by the representations of the type $[\pi, \bar{\tau}]$ of $G_{[m,n]}$ (see (3.8)); note however that although we have

$$G/N_\ell \simeq G_\ell$$

for $\ell \in \Lambda$, this may not extend to arbitrary finite subsets $m \in S(\Lambda)$ (in other words, ρ_m is not necessarily onto, so that G/N_m may be smaller than G_m). The point in both estimates is that they are uniform over all representations that factor through some N_i , and exponentially small as k grows compared to the trivial bounds (namely $\|e\| \|f\|$ or $\dim \pi$, respectively).

Readers unfamiliar with probability theory are invited to translate the proof into the language of words of length N , which is somewhat simpler.

Proof Let $i \in I$ and let π be a representation that factors as

$$\pi : G \rightarrow G_i = G/N_i \rightarrow GL(V),$$

and which has no non-zero invariant vector (i.e., π does not contain the trivial representation of G). Clearly it suffices to prove (7.3) since (7.4) follows, the trace of a matrix being equal to the sum of the diagonal matrix coefficients in an orthonormal basis.

Let

$$M = \mathbf{E}(\pi(\xi_k)) = \sum_{s \in S} p(s)\pi(s), \quad M^+ = \text{Id} - M, \quad M^- = \text{Id} + M,$$

which are elements of the endomorphism ring $\text{End}(V)$, independent of k (since the ξ_k are identically distributed). These elements are self-adjoint because the generating set and the distribution of ξ_k are symmetric, and the representation unitary so that $\pi(s)^* = \pi(s^{-1})$. Further, let

$$N_0 = \mathbf{E}(\pi(X_0)) = \sum_{t \in T} \mathbf{P}(X_0 = t)\pi(t) \in \text{End}(V).$$

The formula $\mathbf{E}(XY) = \mathbf{E}(X)\mathbf{E}(Y)$ when X and Y are independent random variables yields

$$\mathbf{E}(\pi(X_k)) = \mathbf{E}(\pi(X_0)\pi(\xi_1) \cdots \pi(\xi_k)) = N_0 M^k$$

(this is most commonly stated for scalar-valued random variables, but it is easy to check for variables taking values in a matrix ring, see (F.2) in Appendix F), and linearity of the inner product and of the expectation then gives

$$\mathbf{E}(\langle \pi(X_k)e, f \rangle) = \langle \mathbf{E}(\pi(X_k))e, f \rangle = \langle N_0 M^k e, f \rangle = \langle M^k e, N_0^* f \rangle$$

where N_0^* is the adjoint of N_0 .

Let ρ be the spectral radius of M , or equivalently the largest of absolute values of the eigenvalues of M . Note that $0 \leq \rho \leq 1$ since the eigenvalues lie inside the unit disc, by virtue of M being an average of unitary operators. Moreover, since M is self-adjoint, the eigenvalues are real numbers in $[-1, 1]$.

Since N_0 is also an average of unitary operators, so is its adjoint, and hence the norm of N_0^* is at most 1. Hence we have

$$|\langle M^k e, N_0^* f \rangle| \leq \|e\| \|f\| \rho^k,$$

and it only remains to prove the existence of a constant $\delta > 0$, independent of i and π , such that $0 \leq \rho \leq 1 - \delta < 1$; indeed, we may then take $\eta = -\log(1 - \delta) > 0$.

Clearly $\rho = \max(\rho_+, \rho_-)$, where $\rho_+ \in \mathbf{R}$ (respectively ρ_-) is the largest eigenvalue and ρ_- is the opposite of the smallest eigenvalue. We will prove that $\rho_{\pm} \leq 1 - \delta_{\pm}$ with $\delta_{\pm} > 0$ independent of i and π .

For this we use the obvious fact that $1 - \rho^+$ (respectively $1 + \rho^-$) is the smallest eigenvalue of M^+ (respectively M^-). Hence we can use the variational

characterization of the smallest eigenvalue λ of a self-adjoint operator T on a finite-dimensional Hilbert space:

$$\lambda = \min_{v \neq 0} \frac{\langle Tv, v \rangle}{\|v\|^2}$$

(this formula is obvious once T is diagonalized).

So we need to find lower bounds for such quotients when $T = M^\pm$. The crucial facts are the formulas

$$\begin{aligned} \langle M^+ v, v \rangle &= \frac{1}{2} \mathbf{E}(\|\pi(\xi_k)v - v\|^2) = \frac{1}{2} \sum_{s \in S} p(s) \|\pi(s)v - v\|^2, \\ \langle M^- v, v \rangle &= \frac{1}{2} \mathbf{E}(\|\pi(\xi_k)v + v\|^2) = \frac{1}{2} \sum_{s \in S} p(s) \|\pi(s)v + v\|^2, \end{aligned}$$

the proofs of which are identical; for instance

$$\begin{aligned} \mathbf{E}(\|\pi(\xi_k)v - v\|^2) &= \mathbf{E}(\|\pi(\xi_k)v\|^2 - 2 \operatorname{Re}(\langle \pi(\xi_k)v, v \rangle) + \|v\|^2) \\ &= 2\|v\|^2 - 2\langle \mathbf{E}(\pi(\xi_k))v, v \rangle = 2\langle (\operatorname{Id} - M)v, v \rangle \end{aligned}$$

(where we used the fact that $\|\pi(\xi_k)v\| = \|v\|$ because the inner product is G -invariant, and also the fact that M is self-adjoint to dispense with the real part).

Now we find lower bounds for each quotient separately. For ρ_+ , we observe that, quite tautologically, we have

$$\frac{\langle M^+(v), v \rangle}{\|v\|^2} = \frac{1}{2} \sum_{s \in S} p(s) \frac{\|\pi(s)v - v\|^2}{\|v\|^2} \geq \frac{p^+}{2} \inf_{\varpi} \inf_{v \neq 0} \max_{s \in S} \frac{\|\varpi(s)v - v\|^2}{\|v\|^2}, \quad (7.8)$$

where $p^+ = \min p(s) > 0$ by assumption, and where ϖ ranges over *all* unitary representations of G that factor through some N_i and do not contain the trivial representation (and of course $\|\cdot\|$ on the right-hand side is the unitary norm for each such representation). But it is exactly the content of Property (τ) for G with respect to (N_i) that this triple extremum is > 0 (see, e.g., [91, Definition 4.3.1], Appendix D). If we denote it by $\kappa = \kappa(G, S, (N_i))$, this gives the desired inequality with

$$\delta^+ = \frac{\kappa p^+}{2}.$$

Now we come to ρ_- . Here a suitable lower-bound follows from the second assumption of the theorem, by applying Theorem 6.6 of [111] (due to Diaconis, Saloff-Coste, Stroock), using the fact that any eigenvalue of M is also an eigenvalue of M_{reg} , where M_{reg} is the analogue of M for the regular representation of G on $L^2(G/N_i)$.

For completeness, we prove what is needed here, adapting the arguments to the case of a general representation. We need a lower bound for

$$\sum_{s \in S} p(s) \|\pi(s)v + v\|^2,$$

and to find one, we use the word $r = s_1 \cdots s_c$ of odd length c such that r , in G , is trivial. Indeed, for $v \in V$, we can write

$$v = \frac{1}{2} \left((v + \pi(s_1)v) - (\pi(s_1)v + \pi(s_1 s_2 v)) + \cdots + (\pi(s_1 \cdots s_{c-1})v + \pi(1)v) \right)$$

(this is where the odd length is crucial), and hence by Cauchy's inequality we get

$$\|v\|^2 \leq \frac{c}{4} \sum_{i=0}^{c-1} \|\pi(r_i)v + \pi(r_i s_{i+1})v\|^2 = \frac{c}{4} \sum_{i=0}^{c-1} \|v + \pi(s_{i+1})v\|^2$$

(again we use the invariance of the inner product). By positivity, since at worst all s_i are equal to the same generator in S , we get

$$\|v\|^2 \leq \frac{c^2}{4} \sum_{s \in S} \|\pi(s)v + v\|^2 \leq \frac{c^2}{2p^+} \langle M^-(v), v \rangle, \quad (7.9)$$

which, from what we saw, implies that $\rho^- \leq 1 - \delta^-$ with $\delta_- = \frac{2p^+}{c^2} > 0$. \square

One can see in the proof how naturally Property (τ) enters the picture, but of course the tautological lower bound (7.8) might not be best possible, and similar, or slightly weaker, results may well be possible in groups without Property (τ) . Indeed, in the previous chapter, we considered random walks on \mathbf{Z} , which definitely does not satisfy Property (τ) (and, not coincidentally, we obtained a bound with polynomial decay instead of exponential decay).

On the other hand, the second assumption, funny looking as it may seem, is sharp. This is again related to periodicity problems in the random walk. Indeed, if there is a non-trivial homomorphism $G \rightarrow \mathbf{Z}/2\mathbf{Z}$ with the additional condition that $\varepsilon(s) = -1$ for all $s \in S$ (which was the case for $G = \mathbf{Z}$ if $S = \{\pm 1\}$, and may also happen for $SL(2, \mathbf{Z})$ for instance, as we will see later), we can see this map as a representation $\varepsilon : G \rightarrow \{\pm 1\} \subset \mathbf{C}^\times$, for which we have trivially

$$\mathbf{E}(\varepsilon(X_k)) = (-1)^k$$

which shows no cancellation whatsoever as k grows.

It is interesting to interpret the random walks geometrically using the Cayley graph of G with respect to S , and to rephrase the conditions in this manner.

For this purpose, we define formally a graph $\Gamma = (V, E)$ to be the data consisting of a set of vertices V and an edge map $E : V \times V \rightarrow \mathbf{N}$, which

gives the number of edges joining two vertices, with $E(x, y) = E(y, x)$. If $E(x, y) > 0$, we say that x and y are adjacent, or are neighbours. Note that we allow self-loops (if $E(x, x) > 0$) and multiple edges (if $E(x, y) > 1$) in Γ . We can see the vertex set of a graph as a metric space with distance given by $d_\Gamma(x, x) = 0$ and otherwise $d_\Gamma(x, y)$ is the smallest $k \geq 1$ such that there exists a *path* of length k joining x to y , i.e., a sequence $\gamma = (x_0, x_1, \dots, x_k)$ in V^{k+1} with $x_0 = x$, $x_k = y$ and $E(x_i, x_{i+1}) > 0$ for $0 \leq i \leq k - 1$. Using this metric, we can speak in particular of topological or metric properties of the graph, and define related invariants (e.g., connectedness, diameter, etc.).

Here is an example: the graph with six vertices and E determined by the matrix (called the *adjacency matrix*)

$$E = \begin{pmatrix} 0 & 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 2 \\ 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 2 & 0 & 0 \end{pmatrix}$$

is represented in Figure 7.1.

The *Cayley graph* $C_G(H, S)$ associated to a quotient $G \xrightarrow{f} H$ of a finitely generated group G and a system of generators S of G is the graph with $V = H$ and

$$E(g, h) = |\{s \in S \mid gs(s) = h\}|$$

(which may be > 1 if two generators map to the same element in H): from each vertex, there are as many edges exiting as there are elements of S . The graph is connected because S is a generating set of G .

We can see the random walk (X_k) on G as a random walk on $C_G(G, S)$, where the walker, at each step, chooses a neighbour of its position in the graph, which




Figure 7.1 A graph with six vertices

is of the form $X_k s$, and walks there with probability given by $p(s)$. Given a family of finite quotients (N_i) , we have for each i an induced walk on the finite Cayley graph $C_G(N_i, S)$. Then the probabilistic content of Proposition 7.2, for a fixed i , is the well-known fact that the distribution of a random walk of this type on a finite graph converges exponentially fast to the uniform distribution on the set (independently of the distribution of the steps of the walk).

Precisely, such convergence can only occur if the walk is not operating on a bipartite graph. Recall that a graph Γ as above is *bipartite* if V is the disjoint union of two non-empty sets I_Γ and O_Γ , in such a way that $E(x, y) = 0$ if either $\{x, y\} \subset I_\Gamma$ or $\{x, y\} \subset O_\Gamma$ (no edge, including no self-loop, can join two elements in the same part of the graph). Then, clearly, starting from any vertex, a random walk X_k on the graph will always satisfy the condition that X_k is in the same part as x_0 if and only if k is even, and the distribution of X_k is never close to the uniform distribution.

Now coming to our second condition in Proposition 7.2, it may be rephrased as stating that there is in the Cayley graph $C_G(G, S)$ a loop (i.e., a path with identical extremities, starting from the origin) of odd length c . Now it is clear that the existence of such a loop is *equivalent* with the fact that the graph is not bipartite (i.e., that there exists no non-trivial partition of V as union of two sets I_Γ and O_Γ that make it a bipartite graph). Indeed, in a bipartite graph, any path of odd length has its extremities lying in distinct parts, and so cannot be a loop; while if there is no loop of odd length in a connected graph, it can be made into a bipartite graph by fixing a vertex x_0 and defining

$$I_\Gamma = \{x \in V \mid d_\Gamma(x_0, x) \text{ is even}\}, \quad O_\Gamma = \{x \in V \mid d_\Gamma(x_0, x) \text{ is odd}\} \quad (7.10)$$

(if an edge were to go, e.g., from $x \in I_\Gamma$ to $y \in I_\Gamma$, following a path of even length from x_0 to x , then this edge, then coming back to y , would yield a loop of odd length).

In terms of graphs, the proposition is related to the well-known crucial fact that the family of Cayley graphs $C_G(N_i, S)$ for a group G having Property (τ) with respect to a family (N_i) forms an *expanding graph* or *expander family*. We refer to the books by Lubotzky [91] and by Sarnak [112], and to the recent survey by Hoory, Linial and Wigderson [61] for more on expanders and their rather amazing applications (which go well beyond ‘pure’ mathematics).

In many applications, the ‘spectral gap’ (which is the quantity we called ρ^+ during the course of the proof of Proposition 7.2) is emphasized foremost. Indeed, this gap being bounded away from zero is exactly what defines a family of expanding graphs. The issue of ρ^- is much less critical – it can essentially only be too small if the graph is ‘almost’ bipartite, and for many applications this rather irrelevant detail can be disregarded using the simple expedient of

replacing S by $S \cup \{1\}$ (algebraically), or replacing the graph by adding a self-loop on each vertex if none existed; or (probabilistically) by replacing the given walk with a ‘lazier’ version (see (7.2)). Any of these has the effect of making it possible to take the trivial relation and to apply the proposition with $c = 1$. In fact, we can see directly (in the case of the simple random walk with $p(s) = 1/|S|$) what is the effect on M of replacing S by $S' = S \cup \{1\}$ (if $1 \notin S$): we have

$$M_{S'} = \left(1 - \frac{1}{|S'|}\right)M_S + \frac{1}{|S'|},$$

with obvious notation, and so we directly obtain the lower bound

$$\rho_- \geq -1 + \frac{2}{|S'|}.$$

7.3 Applications to arithmetic groups

We now consider concrete instances of the sieve problems discussed in the previous section, leading (in Section 7.5) to the proof of Theorems 1.3 and 1.5.

We work either with the special linear groups $SL(n, \mathbf{Z})$ or with the symplectic groups $Sp(2g, \mathbf{Z})$. For this purpose, we use the notation from algebraic group theory: let \mathbf{G} be either $SL(n)$ or $Sp(2g)$ for some $n \geq 2$ or $g \geq 1$, and for any (commutative unitary) ring A , let $\mathbf{G}(A) = SL(n, A)$ or $Sp(2g, A)$, respectively. Note that if $f : A \rightarrow B$ is a ring homomorphism, there is an induced group homomorphism $\mathbf{G}(A) \rightarrow \mathbf{G}(B)$, and in particular there are reduction maps

$$\mathbf{G}(\mathbf{Z}) \rightarrow \mathbf{G}(\mathbf{F}_\ell)$$

for any prime ℓ .

We begin by stating a few group-theoretical facts which will be useful later on. For $SL(n, \mathbf{Z})$ at least, they are quite well known.

Lemma 7.3 *Let $\mathbf{G} = SL(n)$ or $Sp(2g)$ as above.*

- (1) *The group $\mathbf{G}(\mathbf{Z})$ is finitely generated.*
- (2) *The reduction map $\mathbf{G}(\mathbf{Z}) \rightarrow \mathbf{G}(\mathbf{Z}/m\mathbf{Z})$ is onto for all integers $m \geq 1$.*
- (3) *If $n \geq 3$ or $g \geq 3$, $\mathbf{G}(\mathbf{Z})$ is equal to its commutator subgroup, but not if $n = 2$ or $g = 2$.*
- (4) *Property (τ) holds for the group $G = \mathbf{G}(\mathbf{Z})$ with respect to the family of congruence subgroups $(\text{Ker}(G \rightarrow \mathbf{G}(\mathbf{Z}/d\mathbf{Z})))_{d \geq 1}$.*
- (5) *If $n \geq 3$ or $g \geq 3$, then for any finite symmetric generating set S of $\mathbf{G}(\mathbf{Z})$, there exists a relation of odd length in S .*

Proof

- (1) For $\mathbf{G} = SL(n)$ it is of course well known (by row-and-column reduction of integral matrices to compute elementary divisors) that transvections generate $SL(n, \mathbf{Z})$, and transvections lie in one of the infinite cyclic subgroups generated by an elementary matrix $E_{i,j}$ with 1 on the diagonal and at the (i, j) -th position for $1 \leq i < j \leq n$. Hence the set S of elementary matrices with ± 1 off the diagonal is a finite symmetric generating set of $SL(n, \mathbf{Z})$.

For $\mathbf{G} = Sp(2g)$, an analogue of this is still true, and in fact both cases can be treated in parallel using the theory of algebraic groups. There are finitely many *root subgroups* X_α in \mathbf{G} , isomorphic to the additive group, so that we obtain homomorphisms

$$x_\alpha : \mathbf{Z} \rightarrow \mathbf{G}(\mathbf{Z})$$

and the *elementary subgroup* $E(\mathbf{G}, \mathbf{Z})$ generated by $x_\alpha(1)$ for all α turns out to satisfy

$$E(\mathbf{G}, \mathbf{Z}) = \mathbf{G}(\mathbf{Z})$$

for the groups under consideration (the statement can be found for $Sp(2g)$ in [7, Corollary 12.5] or [54, 5.3.4], though both use slightly different definitions of the elementary subgroup; in both cases, they are still finitely generated). This gives finitely many generators. Concretely, the $x_\alpha(1)$ for $\mathbf{G} = SL(n)$ are precisely the elementary matrices above.

(Even more precisely, explicit *presentations* of $SL(n, \mathbf{Z})$ and $Sp(2g, \mathbf{Z})$ are known: see, e.g., [54, 9.2.13] except for $Sp(4, \mathbf{Z})$, and [8] in this last case.)

- (2) This is proved, e.g., in [120, Lemma 1.38] for the case of $SL(n)$, and in [102, Theorem VII.21] for $Sp(2g)$. Alternatively, one can use the fact that the groups $\mathbf{G}(\mathbf{F}_\ell)$ are also generated by the corresponding root subgroups, i.e., by the images of

$$\tilde{x}_\alpha : \mathbf{F}_\ell \rightarrow \mathbf{G}(\mathbf{F}_\ell),$$

and observe that the generators $x_\alpha(1)$ of $\mathbf{G}(\mathbf{Z})$ reduce to the generators $\tilde{x}_\alpha(1)$ of $\mathbf{G}(\mathbf{F}_\ell)$. Or one could apply the so-called Strong Approximation Theorem, though the latter would be most interesting if dealing with subgroups of $\mathbf{G}(\mathbf{Z})$ where explicit generators are not so easily found . . .

- (3) It suffices to show that elementary matrices are commutators, and this follows for $SL(n)$ from the well-known commutator relations

$$E_{i,j}[E_{i,k}, E_{k,j}]^{-1} = 1$$

if i, j, k are distinct indices $\leq n$, where $E_{i,j}$ is the elementary matrix in (1).

For $Sp(2g)$ with $g \geq 3$, the stated properties can also be obtained by looking at (more complicated) commutator relations among the generators $x_\alpha(1)$ above; the statement itself is a special case of [7, Proposition 13.2], taking into account Corollary 12.5 of [7]; one can also check this using the presentation in [54, 9.2.13].

For $n = 2$, it is well known that $[SL(2, \mathbf{Z}), SL(2, \mathbf{Z})]$ is of index 12 in $SL(2, \mathbf{Z})$, and for $Sp(4, \mathbf{Z})$, that $[Sp(4, \mathbf{Z}), Sp(4, \mathbf{Z})]$ is of index (at least) two, because of (2) and the existence of an exceptional isomorphism

$$Sp(4, \mathbf{F}_2) \simeq \mathfrak{S}_6$$

(see, e.g., [103, 3.1.5]) which gives a non-trivial map $Sp(4, \mathbf{Z}) \rightarrow \mathfrak{S}_6 \xrightarrow{\varepsilon} \{\pm 1\}$. See Section 7.4 for more details on these two cases.

- (4) This crucial fact is well known; in fact, for $n \geq 3$ or $g \geq 2$, the group $\mathbf{G}(\mathbf{Z})$ is a lattice in the group $G(\mathbf{R})$ which is a semisimple algebraic group over \mathbf{R} with \mathbf{R} -rank ≥ 2 (see Appendix E for the terminology, if it is unfamiliar), and hence it satisfies the stronger Property (T) of Kazhdan, which means that in (7.8), the infimum may be taken on *all* unitary representations of G not containing the trivial representation and remains > 0 ; see, e.g., [58, Corollary 3.5], [91, Proposition 3.2.3, Example 3.2.4, Section 4.4]. (Note that this implies (1) again by basic properties of discrete groups with Property (T), see Appendix D.)

For the case of $SL(n, \mathbf{Z})$, $n \geq 3$, we give in Appendix D a fairly complete sketch of the proof of Property (τ), following the approach of Y. Shalom [118], which is quite elementary (the earlier approach of Burger [18] could also be used). For $SL(2)$, see the beginning of Section 7.4.

- (5) Note that if $\mathbf{G} = SL(n)$ and the generating set S is the one of elementary matrices described in (1), the commutator relation stated in (3) (for $(i, j, k) = (1, 2, 3)$, say) is a relation of odd length 5. In the general case, if all S -relations were of even length, the homomorphism

$$\begin{cases} F(S) & \rightarrow & \{\pm 1\} \\ s & \mapsto & -1 \end{cases}$$

(where $F(S)$ is the free group on S) would induce a non-trivial homomorphism $\mathbf{G}(\mathbf{Z}) \rightarrow \{\pm 1\}$. However, there is no such homomorphism for the groups under consideration, since it would have to factor through $\mathbf{G}(\mathbf{Z})/[\mathbf{G}(\mathbf{Z}), \mathbf{G}(\mathbf{Z})] = 1$ by (3). \square

Consider now $G = \mathbf{G}(\mathbf{Z})$, and either the group sieve setting

$$\Psi = (G, \{\text{primes}\}, G \xrightarrow{p\ell} \mathbf{G}(\mathbf{F}_\ell))$$

where the maps ρ_ℓ are simply reduction modulo ℓ (which are onto as we just recalled), or the induced conjugacy sieve setting. We will estimate the large sieve constant arising from a siftable set of the type $\Upsilon = (\Omega, \mathbf{P}, X_k)$ associated to a random walk (X_k) on G as in the previous section.

Theorem 7.4 *Let $\mathbf{G} = SL(n)$, $n \geq 3$, or $Sp(2g)$, $g \geq 3$, be as before, $G = \mathbf{G}(\mathbf{Z})$, and let*

$$\Psi = (G, \{\text{primes}\}, G \rightarrow G_\ell)$$

be the group sieve setting where $G_\ell = \mathbf{G}(\mathbf{F}_\ell)$ for ℓ prime. Let $S = S^{-1}$ be a symmetric generating set for G , and let (X_k) be a symmetric left-invariant random walk on G with identically distributed independent steps (ξ_k) such that

$$\mathbf{P}(\xi_k = s) = \mathbf{P}(\xi_k = s^{-1}) = p(s) > 0, \quad \text{for all } s \in S.$$

- (1) *For any sieve support \mathcal{L} , the large sieve constant for the induced conjugacy sieve satisfies*

$$\Delta(X_k, \mathcal{L}) \leq 1 + R(\mathcal{L}) \exp(-\eta k), \quad (7.11)$$

where $\eta > 0$ is a constant depending only on \mathbf{G} , S and the distribution of (ξ_k) , and¹

$$R(\mathcal{L}) = \max_{m \in \mathcal{L}} \left\{ A_\infty(G_m) \right\} \times \sum_{n \in \mathcal{L}} A_1(G_n).$$

- (2) *For any sieve support \mathcal{L} , the large sieve constant for the group sieve satisfies*

$$\Delta(X_k, \mathcal{L}) \leq 1 + \tilde{R}(\mathcal{L}) \exp(-\eta k), \quad (7.12)$$

where $\eta > 0$ is the same constant as in (1) and

$$\tilde{R}(\mathcal{L}) = \max_{m \in \mathcal{L}} \left\{ \sqrt{A_\infty(G_m)} \right\} \times \sum_{n \in \mathcal{L}} A_{5/2}(G_n)^{5/2}.$$

In terms of words of length N (as in (7.6) and (7.7)), this translates to

$$\Delta(W, \mathcal{L}) \leq |W| + |W|^{1-\alpha} R(\mathcal{L}), \quad \Delta(W, \mathcal{L}) \leq |W| + |W|^{1-\alpha} \tilde{R}(\mathcal{L})$$

in the conjugacy case (respectively the group case), with $\alpha = \eta / \log |S|$.

Proof We first notice that by Lemma 7.3, (4) and (5), both assumptions of Proposition 7.2 hold for G with respect to the family of congruence subgroups $(\text{Ker}(G \rightarrow \mathbf{G}(\mathbf{Z}/d\mathbf{Z})))_{d \geq 1}$.

¹ With notation as in Chapter 5.

- (1) We use the bound arising from duality for a conjugacy sieve, based on the ‘exponential sums’ $W(\pi, \tau)$ of (3.9) for $m, n \in \mathcal{L}$ and $\pi, \tau \in \Pi_m^*, \Pi_n^*$ respectively, namely

$$W(\pi, \tau) = \mathbf{E}(\text{Tr}([\pi, \bar{\tau}] \circ \rho_{[m,n]}(X_k)))$$

in probabilistic notation.

First of all, for any integer $d \geq 1$, we can identify the group G_d , which is defined as the product of G_ℓ for $\ell \mid d$, with $\mathbf{G}(\mathbf{Z}/d\mathbf{Z})$. Indeed, this is simply because of the Chinese Remainder Theorem and the fact that elements of $\mathbf{G}(A)$, for any ring A , are defined by algebraic equations. Since the reduction maps $G \rightarrow \mathbf{G}(\mathbf{Z}/d\mathbf{Z})$ are onto for all d (see Lemma 7.3, (2)), it follows that $\rho_d : G \rightarrow G_d$ is surjective. This applies in particular to $d = [m, n]$ for any squarefree integers m and n .

By Lemma 3.4, the representation $[\pi, \bar{\tau}]$ of $G_{[m,n]}$ defined in (3.3) contains the trivial representation if and only if $(m, \pi) = (n, \tau)$, and then contains it with multiplicity one. Let $[\pi, \bar{\tau}]_0$ denote the orthogonal of the trivial component ($[\pi, \bar{\tau}]_0 = [\pi, \bar{\tau}]$ if $(m, \pi) \neq (n, \tau)$).

We now apply Proposition 7.2 (to G and the family of congruence subgroups) with π replaced by the representation $[\pi, \bar{\tau}]_0 \circ \rho_{[m,n]}$ (which factors through the congruence subgroup $\text{Ker}(G \rightarrow \mathbf{G}(\mathbf{Z}/[m, n]\mathbf{Z}))$). By (7.7), we have

$$|\mathbf{E}(\text{Tr}([\pi, \bar{\tau}]_0 \circ \rho_{[m,n]}(X_k)))| \leq (\dim \pi)(\dim \tau) \exp(-\eta k)$$

where $\eta > 0$ is given by (7.5); note that it depends only on G , S and the distribution of the steps (ξ_k) . Since

$$\text{Tr}([\pi, \bar{\tau}] \circ \rho_{[m,n]}(X_k)) = \delta(\pi, \tau) + \text{Tr}([\pi, \bar{\tau}]_0 \circ \rho_{[m,n]}(X_k))$$

(because the G -invariant subspace has dimension $\delta(\pi, \tau)$), it follows that

$$\left| W(\pi, \tau) - \delta(\pi, \tau) \right| \leq (\dim \pi)(\dim \tau) \exp(\eta k),$$

and from Proposition 2.9, we obtain immediately

$$\Delta(X_k, \mathcal{L}) \leq 1 + \exp(-\eta k) \max_{m \in \mathcal{L}} A_\infty(G_m) \sum_{n \in \mathcal{L}} A_1(G_n),$$

as stated.

- (2) The argument is similar, except that now we use the basis of matrix coefficients for the group sieve setting, and correspondingly we appeal to (7.6) and the fact (see the final paragraphs of Chapter 3 and Proposition 3.6) that the sums $W(\varphi_{\pi, e, f}, \varphi_{\tau, e', f'})$ are (up to the factor $\sqrt{(\dim \pi)(\dim \tau)}$) that

appears in their definition) of the type considered in (7.6), namely they are

$$\mathbf{E} \left(\langle [\pi, \bar{\tau}](\rho_{[m,n]}(X_k))e, f \rangle \right)$$

for some vectors e and f in the space of $[\pi, \bar{\tau}]$.

Before applying Proposition 7.2, we must again isolate the contribution of the trivial representation. Now, $[\pi, \bar{\tau}]$, as stated before, has invariant vectors if and only if $(m, \pi) = (n, \tau)$. Also, we note that if π acts on V_π , then the representation $[\pi, \bar{\tau}] = \pi \otimes \bar{\pi}$ of G_m is isomorphic with the representation on $V_\pi \otimes V'_\pi \simeq \text{End}(V_\pi)$ given by

$$(g, A) \mapsto \pi(g)A\pi(g)^{-1} \quad \text{for } A \in \text{End}(V_\pi).$$

In this description, the space of invariant vectors is one-dimensional in $\text{End}(V_\pi)$, and is spanned by scalar multiples of the identity (which are clearly invariant!), and the orthogonal projection of a linear map $A \in \text{End}(V_\pi)$ onto this space is the scalar multiplication by $\text{Tr}(A)/\sqrt{(\dim \pi)}$ (this is a corollary of the orthogonality relations; note that $\|\text{Id}\|^2 = \dim \pi$, so a normalized generator of the space of homotheties is multiplication by $(\dim \pi)^{-1/2}$).

Now apply this to a rank 1 linear map of the form

$$A = e \otimes e' : v \mapsto \langle v, e' \rangle e$$

where $e, e' \in V_\pi$; the projection to the invariant subspace of this map is the multiplication by

$$\frac{\text{Tr}(A)}{\sqrt{\dim \pi}} = \frac{\langle e, e' \rangle}{\sqrt{\dim \pi}}.$$

This means that for any $g \in G_m$, we have

$$(\pi \otimes \bar{\pi})(g)(e \otimes e') = \frac{\langle e, e' \rangle}{\sqrt{\dim \pi}} + [\pi, \bar{\pi}]_0(g)(e \otimes e')$$

with $[\pi, \bar{\pi}]_0$ as before; applying the scalar product (in $\text{End}(V_\pi)$) with another rank 1 map $f \otimes f'$, we find

$$\langle (\pi \otimes \bar{\pi})(g)(e \otimes e'), f \otimes f' \rangle = \frac{\langle e, e' \rangle \langle f, f' \rangle}{\dim \pi} + \langle [\pi, \bar{\pi}]_0(g)(e \otimes e'), f \otimes f' \rangle.$$

The point is now that in our situation, e, e' (respectively f, f') are all taken from a fixed orthonormal basis of V_π , and hence the leading term is zero except when $(e, f) = (e', f')$, in which case taking the expectation contributes $1/\dim \pi$.

Coming back to the general case, the contribution of $[\pi, \tau]_0$ is always handled by (7.6), and we derive

$$\left| W(\varphi_{\pi, e, f}, \varphi_{\tau, e', f'}) - \delta((\pi, e, f), (\tau, e', f')) \right| \leq \sqrt{(\dim \pi)(\dim \tau)} \exp(-k\eta)$$

with the same value of η as before, and hence

$$\begin{aligned} \Delta(X_k, \mathcal{L}) &\leq 1 + \exp(-k\eta) \max_{m, \pi, e, f} \sqrt{\dim \pi} \sum_{n \in \mathcal{L}} \sum_{\tau, e', f'} \sqrt{\dim \tau} \\ &\leq 1 + \exp(-k\eta) \max_{m \in \mathcal{L}} \sqrt{A_\infty(G_m)} \sum_{n \in \mathcal{L}} \sum_{\tau} (\dim \tau)^{5/2} \\ &= 1 + \exp(-k\eta) \max_{m \in \mathcal{L}} \sqrt{A_\infty(G_m)} \sum_{n \in \mathcal{L}} A_{5/2}(G_n)^{5/2} \end{aligned}$$

which is the estimate we claimed. \square

Remark 7.5 In applications, this means that \mathcal{L} may be chosen at will, provided that $R(\mathcal{L})$ (or $\tilde{R}(\mathcal{L})$) is somewhat smaller than $\exp(\eta k)$. The sharpest estimates for $R(\mathcal{L})$ require bounds such as those proved for finite groups of Lie type in Chapter 5 (see Proposition 5.4); however there is no point in applying those fairly sophisticated results if no explicit value of η is known, since the ‘trivial’ bounds of Proposition 5.2 are qualitatively equivalent.

From (7.5), we see that computing η requires knowing an explicit value of the (T) -constant (or (τ) -constant with respect to the congruence subgroups, in this case) for G . The question of such explicit bounds was first raised by Serre, de la Harpe and Valette, and we see that this is clearly a natural question with concrete applications, such as explicit sieve bounds (other important applications, already well established, are explicit expander bounds, though Ramanujan graphs, which are the best expanders, are not of this type). In Section 7.7, we will describe an impressive example due to Shalom (and Kassabov) and its use for sieve.

7.4 The cases of $SL(2)$ and $Sp(4)$

The group-theoretic Lemma 7.3 has shown that the ‘small rank’ groups $\mathbf{G}(\mathbf{Z})$ for $\mathbf{G} = SL(2)$ or $Sp(4)$ need to be treated separately. Indeed, the existence of non-trivial homomorphisms to $\{\pm 1\}$ factoring through $\mathbf{G}(\mathbf{F}_2)$ indicates that equidistribution modulo 2 does not hold.

However, provided the sieve applications can be dealt with using only odd primes (or primes ≥ 5 for the case of $SL(2)$), it is possible to still derive fairly

good results, as we will describe. In fact, we use two different techniques, one for $SL(2, \mathbf{Z})$ which gives weaker results, and another for $Sp(4, \mathbf{Z})$ which essentially recovers sieve bounds of the same quality as in the previous section. The two methods could be exchanged, and we could develop an analogue of the more efficient one for $SL(2, \mathbf{Z})$, but we refrain from doing so simply to illustrate the possibilities available; for some other applications, it is possible that the technique used for $Sp(4, \mathbf{Z})$ does not work easily.

We start with $G = SL(2, \mathbf{Z})$. The reason sieve is still possible is of course that, although G does not have Property (T), it remains true that G has Property (τ) with respect to the family of congruence subgroups $\Gamma(d) = \text{Ker}(SL(2, \mathbf{Z}) \rightarrow SL(2, \mathbf{Z}/d\mathbf{Z}))$, and this is the main ingredient we need for our basic sieve. This last result is in fact deeper than Property (T) for $SL(n, \mathbf{Z})$, $n \geq 3$. It comes, in the final analysis, from Selberg's theorem that the smallest positive eigenvalue of the hyperbolic Laplacian acting on square-integrable functions on the quotient $\Gamma(d) \backslash \mathbf{H}$ satisfies $\lambda_1 \geq 3/16$ for all $d \geq 1$. See, e.g., [66, Theorem 11.6] for a proof of this result, noting that any bound $\lambda_1 \geq c > 0$ for all d would be qualitatively sufficient, and that a famous conjecture of Selberg states that, in fact, $\lambda_1 \geq 1/4$, which is optimal. The link with Property (τ) is not obvious, and is explained in [91, Theorem 4.3.2, (vi) implies (i)], for instance, though it is stated only in the case of eigenvalues of compact Riemann surfaces, whereas the quotients $\Gamma(d) \backslash \mathbf{H}$ are finite volume non-compact hyperbolic surfaces. The extension of the result to this situation is due to Brooks [16, Corollary p. 182].

The real problem in extending the sieve to $SL(2, \mathbf{Z})$ is the periodicity constraint on the random walk: the existence of relations of odd length is not true for all generating sets S . For instance, consider the homomorphism

$$SL(2, \mathbf{Z}) \rightarrow SL(2, \mathbf{F}_2) \simeq \mathfrak{S}_3 \xrightarrow{\varepsilon} \{\pm 1\},$$

where the isomorphism in the middle is obtained by looking at the action on the three lines in \mathbf{F}_2^2 , and ε is the signature. All four elements of the symmetric generating set

$$S = \left\{ \begin{pmatrix} 1 & \pm 1 \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ \pm 1 & 1 \end{pmatrix} \right\}, \quad (7.13)$$

of $SL(2, \mathbf{Z})$ map to transpositions in \mathfrak{S}_3 , so we have $\varepsilon(r) = -1$ for any word of odd length in the alphabet S .

However, although this proves that Proposition 7.2 can not be applied to general random walks with arbitrary generating sets for the family of all congruence subgroups, there are a number of ways to obtain similar sieve results:

- One may simply assume that $1 \in S$, adjoining it to S if need be (while changing the probabilities $p(s)$, e.g., $\mathbf{P}(\xi_k = s) = \frac{1}{2}p(s)$ and $\mathbf{P}(\xi_k = 1) = \frac{1}{2}$,

which preserves the relative probabilities of the non-trivial steps). No other change is needed in Proposition 7.2, and Theorem 7.4 holds under this condition. Such an approach is the most natural when the random walk is thought of as a *tool* for some other purpose: there is no reason not to set aside in this manner the (irrelevant) issues of periodicity.

- More generally, one may consider only generating sets that do satisfy the desired condition.
- Finally, in order to investigate a random walk for itself (if only as a challenge), one may observe that it remains true that for an arbitrary symmetric set of generators S of $SL(2, \mathbf{Z})$, and for any integer n coprime with 6, there is a relation of odd length in $SL(2, \mathbf{Z}/n\mathbf{Z})$ with respect to the reductions modulo n of the set S , simply because $SL(2, \mathbf{Z}/n\mathbf{Z})$ has no non-trivial homomorphism $SL(2, \mathbf{Z}/n\mathbf{Z}) \rightarrow \{\pm 1\}$ under these conditions (indeed, $SL(2, \mathbf{Z}/n\mathbf{Z})$ is then equal to its commutator subgroup; this follows from the case of prime power order, which itself is reduced to the case of finite fields \mathbf{F}_ℓ with $\ell > 3$, for which this is well known, see, e.g., [86, Theorem 8.3]).

In terms of the Cayley graph, this means that $C_{SL(2, \mathbf{Z})}(SL(2, \mathbf{Z}/n\mathbf{Z}, S))$ contains some cycle of odd length c_n . We can use this in Proposition 7.2, with the proviso that the value of η is not constant anymore, but may depend on n ; precisely by (7.5), we have

$$\eta \geq p^+ \min \left(\frac{2}{c_n^2}, \frac{\kappa}{2} \right) \gg \frac{p^+}{c_n^2} \quad (7.14)$$

for $n \geq 1$, the implied constant depending on S .

The size of c_n may be estimated using the following general upper result for testing whether a graph is bipartite:

Lemma 7.6 *Let $\Gamma = (V, E)$ be a finite graph. Then if Γ is not bipartite, there exists in Γ a cycle of odd length $c(\Gamma) \leq 2\delta + 1$, where δ is the diameter of Γ .*

Proof Fix some vertex $x_0 \in \Gamma$ and consider two vertices y and z which are neighbours but satisfy $d(x_0, y) \equiv d(x_0, z) \pmod{2}$; these exist, as we already observed at the end of Section 7.2, because otherwise the graph would be bipartite, with I_Γ and O_Γ given by (7.10).

Note that we have $d(x, y) = d(x, z)$, because we know that $|d(x, y) - d(x, z)| \leq 1$ anyway (y and z are adjacent). Now, following a path γ_1 of length $d(x, y) \leq \delta$ from x to y , then the edge from y to z , then a path of length $d(z, x) = d(x, y)$ from z to x , we obtain a loop in Γ of odd length $2d(x, y) + 1 \leq 2\delta + 1$. \square

Remark 7.7 The example of a cycle of odd length, i.e., of the Cayley graph of $\mathbf{Z}/m\mathbf{Z}$ with respect to $S = \{\pm 1\}$ for odd $m \geq 3$, shows that this is best possible for arbitrary graphs.

Moreover, in our case, the order of magnitude can not be improved, as follows from the fact that the *girth* of the Cayley graph (i.e., the length of the shortest relation which is non-trivial in the free group F_S , whether of even or odd length), is $\gg \log n$ (at least when n is prime, see, e.g., [47]).

From this we can now prove:

Proposition 7.8 *Let $G = SL(2, \mathbf{Z})$ and consider the group sieve setting*

$$\Psi = (G, \{\text{primes}\}, G \rightarrow SL(2, \mathbf{F}_\ell))$$

and its associated conjugacy sieve, with the siftable set associated to an arbitrary symmetric left-invariant random walk (X_k) on G with respect to a finite symmetric generating set S .

(1) *If $1 \in S$, then for any sieve support \mathcal{L} , we have*

$$\Delta(X_k, \mathcal{L}) \leq 1 + \tilde{R}(\mathcal{L}) \exp(-k\eta), \quad \text{for the group sieve}$$

$$\Delta(X_k, \mathcal{L}) \leq 1 + R(\mathcal{L}) \exp(-k\eta), \quad \text{for the conjugacy sieve}$$

for all $k \geq 1$, where $\eta > 0$ depends only on S and the distribution of the steps of the random walk, and where $R(\mathcal{L})$ and $\tilde{R}(\mathcal{L})$ are the same as in Theorem 7.4, applied for G .

(2) *If S is an arbitrary generating set, then for any prime sieve support \mathcal{L}^* not containing 2 and 3 and for any associated sieve support \mathcal{L} such that $\max \mathcal{L} \leq L$, $L \geq 6$, we have*

$$\Delta(X_k, \mathcal{L}) \leq 1 + \tilde{R}(\mathcal{L}) \exp\left(-\frac{ck}{(\log L)^2}\right), \quad \text{for the group sieve}$$

$$\Delta(X_k, \mathcal{L}) \leq 1 + R(\mathcal{L}) \exp\left(-\frac{ck}{(\log L)^2}\right), \quad \text{for the conjugacy sieve,}$$

for all $k \geq 1$, for some constant $c > 0$ depending only on S and the distribution of the steps of the random walk.

For applications, the exact values of A_p for $SL(2, \mathbf{Z}/m\mathbf{Z})$ for m squarefree are given in (5.16); in particular

$$R(\mathcal{L}) \leq \left(\max_{m \in \mathcal{L}} \psi(m) \right) \sum_{n \in \mathcal{L}} n \psi(n).$$

Proof We have already explained why (1) holds, so we look at the second part. Leaving the case of the group sieve to the reader, in the conjugacy case we

use the adaptation of Proposition 7.2 sketched above and the same argument as in the beginning of the proof in Theorem 7.4 to find that the exponential sums $W(\pi, \tau)$ satisfy

$$|W(\pi, \tau) - \delta(\pi, \tau)| \leq (\dim \pi)(\dim \tau) \exp(-k\eta_{[m,n]})$$

where

$$\eta_{[m,n]} \gg \frac{2p^+}{(2\delta([m,n] + 1)^2)}$$

by the above lemma and (7.14), the implied constant depending on S . Now, the diameter $\delta(d)$ of the Cayley graph of $SL(2, \mathbf{Z}/d\mathbf{Z})$ with respect to the images of the generators of S satisfies

$$\delta(d) \ll \log d$$

for all $d \geq 2$, the implied constant depending only on d . This is indeed a well-known consequence of the definition of expanding graphs (see, e.g., [61, 2.4, p. 17]); roughly speaking, the size of balls in the Cayley graphs are uniformly exponentially increasing until they contain at least half of the vertices, and two such balls (of radius at most logarithmic in the size of the graph) around two points x and y must intersect, so that the diameter is at most twice that radius.

By the definition of L , this means that for some constant $c > 0$ depending only on S , we have

$$|W(\pi, \tau) - \delta(\pi, \tau)| \leq (\dim \pi)(\dim \tau) \exp\left(-\frac{ck}{(\log L)^2}\right)$$

and the estimate of the large sieve constant follows as usual. \square

Remark 7.9 If we take the generating set S of (7.13), the reader is invited to check that the commutator relation

$$\left[\begin{pmatrix} a & 0 \\ 0 & a^{-1} \end{pmatrix}, \begin{pmatrix} 1 & b \\ 0 & 1 \end{pmatrix} \right] = \begin{pmatrix} 1 & b(a^2 - 1) \\ 0 & 1 \end{pmatrix}$$

(usually used in proving that $PSL(2, \mathbf{F}_q)$ is simple for $q > 3$; see, e.g., [86, Theorem 8.3]), although it *does* yield examples of relations with odd length in $SL(2, \mathbf{F}_\ell)$ for $\ell > 3$, does not give in any obvious way a relation of short length. The issue is to write the diagonal matrix on the left-hand side, with $a^2 \neq 1$, as a short product of elements of S . The standard expression

$$\begin{pmatrix} a & 0 \\ 0 & a^{-1} \end{pmatrix} = \begin{pmatrix} 1 & a \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -a^{-1} & 1 \end{pmatrix} \begin{pmatrix} 1 & -1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$$

only gives a short word if both a and a^{-1} have small representatives in \mathbf{Z} ; but this is impossible if $a \in \mathbf{F}_\ell - \{0, \pm 1\}$ if ‘small’ means of logarithmic size

with respect to ℓ . The expanding property of $SL(2, \mathbf{F}_\ell)$ is truly a deep fact. (See [87] for the best current – probabilistic – *algorithm* to express an element in $SL(2, \mathbf{F}_\ell)$ as a short product of the generators (7.13).)

We now describe a way to deal with the situation of $G = Sp(4, \mathbf{Z})$. If S is a symmetric generating set of G for which there exists a relation of odd length, we can simply apply the results of the previous section, so we will work under the opposite assumption. First of all, here are the precise group-theoretic facts which are relevant:

Lemma 7.10 *Let $\mathbf{G} = Sp(4)$, $G = \mathbf{G}(\mathbf{Z}) = Sp(4, \mathbf{Z})$ and let $G' = [G, G]$ be the commutator subgroup. Then there exists a surjective homomorphism $G \xrightarrow{\pi} \mathfrak{S}_6$ such that:*

- (1) *We have $G' = \pi^{-1}(A_6)$, hence G' is of index 2 in G .*
- (2) *The group G' is finitely generated and for a symmetric generating set S of G such that no relation of odd length exists, the finite set $S^2 = S \cdot S$ is a symmetric generating set of G' .*
- (3) *For every odd prime ℓ , the reduction map*

$$G' \rightarrow Sp(4, \mathbf{F}_\ell)$$

is surjective.

- (4) *Property (T) holds for G' .*

Proof The existence of the homomorphism π has already been noted in the proof of Lemma 7.3: it results from the composition of the surjective homomorphisms

$$G \rightarrow \mathbf{G}(\mathbf{F}_2) = Sp(4, \mathbf{F}_2) \simeq \mathfrak{S}_6.$$

For (1), the inclusion $G' \subset \pi^{-1}(A_6)$ is obvious. For the reverse, the following argument may be too complicated, but it works: from [7, Proposition 13.2],² we see that $G' \supset \text{Ker}(G \rightarrow \mathbf{G}(\mathbf{F}_2))$ (the congruence subgroup of level 2, which is the group denoted $Sp_4(A, q)$ in [7] with $A = \mathbf{Z}$ and $q = 2\mathbf{Z}$) so G' is determined by its image in $\mathbf{G}(\mathbf{F}_2)$. Since G maps onto $\mathbf{G}(\mathbf{F}_2)$, the reduction of G' is the derived group of $Sp(4, \mathbf{F}_2)$, hence corresponds to $A_6 \subset \mathfrak{S}_6$. An alternative argument, not necessarily simpler, is to use the presentation in [8] and compute the elementary divisors of the matrix corresponding to the abelianized relations defining $Sp(4, \mathbf{Z})$, to check that the abelian group $G/[G, G]$ is of order 2 (compare with the proof of Proposition 6.4).

² Note that the paper of Bass referenced in [7] for the proof of this proposition never appeared; however, Bass completes the proof in ‘Unitary algebraic K -theory’, in *Lecture Notes in Mathematics* vol. 343, ed. H. Bass (Springer, 1973), p. 257.

Part (2) is now immediate, since the absence of relations of odd length means that all elements of S map to -1 under the composition of this map and the signature $\mathfrak{S}_6 \rightarrow \{\pm 1\}$, hence any product of two elements of S lies in $\pi^{-1}(A_6) = G'$.

Part (3) is a consequence of the surjectivity of $G \rightarrow Sp(4, \mathbf{F}_\ell)$ and the standard fact that the latter group is its own commutator subgroup for ℓ odd.

Part (4), finally, is because G itself has Property (T) and this passes to any finite index subgroup (see Appendix D). \square

As one can guess, we will reduce our sieve problems to sieves on G' using the generating set S^2 ; note that $1 \in S^2$, but this is not cheating, as we will see. Indeed, it is possible to show that $[G', G'] = G'$, so even after removing 1 from S^2 , there exist relations of odd length. However, in reducing a sieve on G to one on G' , it is really S^2 itself which will occur.

Proposition 7.11 *Let $G = Sp(4, \mathbf{Z})$, $G' = [G, G]$, and consider the group sieve setting*

$$\Psi = (G, \{\text{primes}\}, G \rightarrow Sp(4, \mathbf{F}_\ell))$$

and its associated conjugacy sieve, with the siftable set associated to an arbitrary symmetric left-invariant random walk (X_k) on G with respect to a finite symmetric generating set S .

For any prime sieve support \mathcal{L}^ and for any associated sieve support \mathcal{L} , with the only restriction that all elements of \mathcal{L}^* are odd if no element of S lies in G' , we have*

$$\Delta(X_k, \mathcal{L}) \leq 1 + \tilde{R}(\mathcal{L}) \exp(-k\eta), \quad \text{for the group sieve}$$

$$\Delta(X_k, \mathcal{L}) \leq 1 + R(\mathcal{L}) \exp(-k\eta), \quad \text{for the conjugacy sieve}$$

for all $k \geq 1$, where $\eta > 0$ depends only on S and the distribution of the steps of the random walk, and where $R(\mathcal{L})$ and $\tilde{R}(\mathcal{L})$ are the same as in Theorem 7.4, applied for G .

Proof Only the case where no element of S lies in G' requires proof, since otherwise there exists a relation of odd length and we can argue as in Theorem 7.4. We denote by ε the surjective map $G \rightarrow \{\pm 1\}$ so that $G' = \text{Ker } \varepsilon$, and we write $G'_- = \varepsilon^{-1}(-1) = G - G'$ the other coset.

We will use two auxiliary siftable sets, associated to random walks on G' with respect to the generating set $S^2 = S \cdot S \subset G'$. The first is defined by $Y_k = X_{2k}$ for $k \geq 0$, so $Y_0 = 1$ and $Y_{k+1} = Y_k \omega_{k+1}$ with steps given by

$$\omega_k = \xi_{2k-1} \xi_{2k}$$

which is (almost surely) in S^2 . The independence of the original steps implies that (ω_k) is also a sequence of independent S^2 -valued random variables, with symmetric distribution given by

$$\mathbf{P}(\omega_k = t) = \sum_{\substack{s_1, s_2 \in S \\ s_1 s_2 = t}} p(s) > 0$$

for $t \in S^2$. We then argue with the group sieve setting

$$(G', \{\text{odd primes}\}, G' \rightarrow Sp(4, \mathbf{F}_\ell))$$

or the corresponding conjugacy sieve, and the siftable set

$$(\Omega, \mathbf{P}, Y_k)$$

in the same manner as in Theorem 7.4, obtaining

$$\Delta(Y_k, \mathcal{L}) \leq 1 + \tilde{R}(\mathcal{L}) \exp(-k\eta), \quad \text{or } 1 + R(\mathcal{L}) \exp(-k\eta)$$

by means of Property (T) for G' (note that the representations which occur, coming from integers divisible by odd primes, are the same for G or G' , so the quantities $R(\mathcal{L})$ and $\tilde{R}(\mathcal{L})$ are unchanged). For any k , we have tautologically

$$\Delta(X_{2k}, \mathcal{L}) = \Delta(Y_k, \mathcal{L})$$

hence the result for all even steps of the original random walk.

The odd steps, which lie in the second coset G'_- are handled (naturally enough) using the coset sieve setting associated with the exact sequence

$$1 \rightarrow G' \rightarrow G \rightarrow \{\pm 1\} \rightarrow 1$$

and $\alpha = -1$. This is a very easy case because, for any odd prime ℓ , we have $Y_\ell = \rho_\ell(G'_-)^\# = \mathbf{G}(\mathbf{F}_\ell)^\#$, corresponding to the fact that $G' \rightarrow \mathbf{G}(\mathbf{F}_\ell)$ is onto (indeed, if g_0 is an arbitrary fixed element of G'_- , and $y \in \mathbf{G}(\mathbf{F}_\ell)$, we can find an element in G' mapping to $\rho_\ell(g_0^{-1})y$, and then $g_0 y \in G'_-$ maps to y). From this, it follows that the representations and basis functions for this coset sieve setting are the same as those for the sieve setting on G' .

To reduce to the framework of random walk on a group, defining the siftable set requires writing³

$$X_{2k+1} = \xi_1 v_1 \cdots v_k$$

where $v_k = \xi_{2k} \xi_{2k+1}$, which is an S^2 -valued random variable, and (v_k) is a sequence of independent random variables with the same distribution as the sequence (ω_k) used above.

³ Instead of using this trick, one could of course develop directly the appropriate random walks on cosets.

Then we have $X_{2k+1} = Z_k$, where Z_k is a random walk on G' with steps identically distributed as the steps of (Y_k) , but with an initial distribution $Z_0 = \xi_1$ which is random. However, this starting point is independent of the steps (ν_k) and supported on a finite set, so estimating the relevant exponential sums

$$\mathbf{E}(\mathrm{Tr}(\pi(Z_k)))$$

enters into the general context of Proposition 7.2. It is then clear that we obtain the desired result as in the proof of Theorem 7.4, and we leave the details to the reader. \square

7.5 Arithmetic applications

With Theorem 7.4 in hand, it is now easy to prove some concrete results leading to Theorems 1.3 and 1.5. We emphasize once more that this is but one illustration of the sieve.

Theorem 7.12 *Let $\mathbf{G} = SL(n)$, $n \geq 3$, or $Sp(2g)$, $g \geq 2$, be as before, $G = \mathbf{G}(\mathbf{Z})$, and let*

$$\Psi = (G, \{\text{primes}\}, G \rightarrow G_\ell)$$

be the group sieve setting. Let $S = S^{-1}$ be a symmetric generating set for G , and let (X_k) be a symmetric left-invariant random walk on G with identically distributed independent steps (ξ_k) such that

$$\mathbf{P}(\xi_k = s) = \mathbf{P}(\xi_k = s^{-1}) = p(s) > 0, \quad \text{for all } s \in S.$$

Moreover, let W be the Weyl group of \mathbf{G} , i.e., $W = \mathfrak{S}_n$ if $\mathbf{G} = SL(n)$, or $W = W_{2g}$, the group of signed permutations of g pairs of elements if $\mathbf{G} = Sp(2g)$.

(1) *There exists $\eta > 0$ such that*

$$\mathbf{P}(\det(T - X_k) \text{ has Galois group not isomorphic to } W) \ll \exp(-k\eta), \quad (7.15)$$

where η and the implied constant depend only on \mathbf{G} and S . In particular,

$$\mathbf{P}(\det(T - X_k) \in \mathbf{Z}[T] \text{ is reducible}) \ll \exp(-k\eta).$$

(2) *There exists $\beta > 0$ such that*

$$\mathbf{P}(\text{one entry of } X_k \text{ is a square of an integer}) \ll \exp(-k\beta), \quad (7.16)$$

where β and the implied constant depend only on \mathbf{G} and S .

In terms of random products of length N , we have for instance

$$|\{w \in W \mid \det(T - F_w) \in \mathbf{Z}[T] \text{ is reducible}\}| \ll |W|^{1-\alpha},$$

where $\alpha = \eta / \log |S|$, and similarly for the second result.

Proof

- (1) The first thing to do is to explain why the condition in (7.15) is a natural one. For $SL(n, \mathbf{Z})$, this is clear, since a polynomial of degree n has a splitting field with Galois group isomorphic to a subgroup of $\mathfrak{S}_n = W$. For $g \in Sp(2g, \mathbf{Z})$, we have the self-reciprocity property of $P(T) = \det(T - g)$:

$$T^{2g} P(1/T) = P(T) \quad (7.17)$$

which implies that whenever α is a root of P , so is α^{-1} , and hence⁴ that we can arrange the roots of P in \mathbf{C} in g pairs (α, α^{-1}) which are preserved by the Galois group of P . In other words, this Galois group is naturally isomorphic to a subgroup of $W_{2g} = W$, and the statement we want to prove is that, in the overwhelming majority of cases, it is not a proper subgroup. All this is also discussed in detail, in a similar context, in Section 8.1; and in Proposition E.1 in Appendix E, we give a more conceptual proof of this fact.

To obtain (7.15), we apply the large sieve inequality for group sieves of Proposition 3.5, using (7.11). This follows the same strategy as the proof of Gallagher's Theorem (Theorem 4.2, and its adaptation to 'self-reciprocal polynomials', see [27]), but here we do not seek uniformity in terms of the degree (i.e., in terms of the size of the matrices), because usually the Kazhdan or (τ) -constant, which occur in the large sieve constant when applying Theorem 7.4, are not effective. However, in Section 7.7, we will describe a special case where it is possible to have explicit forms of the sieve using (T) -constants for $SL(n, \mathbf{Z})$, due to Shalom and Kassabov. Note also that in Chapter 8, we will refine this type of argument even further, in another context.

We select the prime sieve support $\mathcal{L}^* = \{\ell \leq L\}$ for some $L \geq 2$, and take $\mathcal{L} = \mathcal{L}^*$ (pedantically, the singletons of elements of \mathcal{L}^* , as usual).

Let $d = \dim \mathbf{G}$, which is either $n^2 - 1$ for $SL(n)$ or $2g^2 + g$ for $Sp(2g)$. By Theorem 7.4, we have

$$\Delta(X_k, \mathcal{L}) \leq 1 + R(\mathcal{L}) \exp(-k\eta)$$

⁴ Looking a bit carefully at the case when there are roots ± 1 .

for some $\eta > 0$ depending only on \mathbf{G} and S , and moreover a fairly crude bound gives

$$R(\mathcal{L}) \ll L^{3d/2+1} \quad (7.18)$$

for $L \geq 2$, the implied constant depending on \mathbf{G} (simply by applying Proposition 5.2 with $|\mathbf{G}(\mathbf{F}_\ell)| \ll \ell^d$).

Applying Gallagher's strategy (see the proof of Theorem 4.2), we fix a conjugacy class c in W and proceed to estimate the probability that the Galois group of $\det(T - X_k)$ does not contain an element in the conjugacy class c when seen as subgroup of W .

Thus, the sieving sets are defined as the sets $\Omega_{c,\ell} \subset \mathbf{G}(\mathbf{F}_\ell)$ of matrices with characteristic polynomial in $\mathbf{F}_\ell[T]$ which factor according to the conjugacy class c . From Lemma B.2 and Lemma B.5, it follows that

$$\frac{|\Omega_{c,\ell}|}{|G_\ell|} \gg 1 \quad (7.19)$$

for $L > 2$ where the implied constant depends on \mathbf{G} and c (this asymptotic lower bound is much cruder than what is actually proved in the Appendix, which is uniform in terms of n and g).

Using the fact that having small splitting field means that some conjugacy class does not occur in the Galois group, we have

$$\begin{aligned} & \mathbf{P}(\det(T - X_k) \text{ has Galois group not isomorphic to } W) \\ & \leq \sum \mathbf{P}(S(X, \Omega_c; \mathcal{L}^*)), \end{aligned}$$

and by the large sieve we have

$$\mathbf{P}(S(X, \Omega_c; \mathcal{L}^*)) \leq \Delta(X_k, \mathcal{L}) H^{-1} \ll (1 + L^{3d/2+1} \exp(-k\eta)) H^{-1}$$

by Proposition 3.5, the implied constant depending on \mathbf{G} , c and S , with

$$H = \sum_{\ell \leq L} \frac{|\Omega_{c,\ell}|}{|G_\ell|} \gg \pi(L) \gg \frac{L}{\log L}$$

for $L > 2$ by (7.19). The parameter L may now be chosen to be $L = \exp(2k\eta/(3d+2))$, if this quantity is ≥ 2 , giving

$$\mathbf{P}(\det(T - X_k) \text{ has small Galois group}) \ll (\log L) L^{-1} \ll \exp(-k\eta')$$

where $\eta' > 0$ is any positive real number smaller than $2\eta/(3d+2)$. To account for those k for which $\exp(2k\eta/(3d+2)) < 2$, we need only adjust the value of η' or increase the implied constant.

- (2) Clearly, it suffices to prove the estimate (7.16) for the probability that the (i, j) -th component of X_k is a square, where i and j are fixed integers (from 1 to n or $2g$ in the $SL(n)$ and $Sp(2g)$ cases respectively).

Since the stated condition is not invariant under conjugation, we use the group sieve and use Theorem 7.4, (2), to estimate the large sieve constant for the sieve where $\mathcal{L}^* = \{\ell \leq L\}$, $\mathcal{L} = \mathcal{L}^*$:

$$\Delta(X_k, \mathcal{L}) \leq 1 + \tilde{R}(\mathcal{L}) \exp(-k\eta)$$

for some $\eta > 0$ depending only on \mathbf{G} and S . From Proposition 5.2, we derive the easy bound

$$\tilde{R}(\mathcal{L}) \ll L^{7d/4+1}$$

for $L \geq 2$, where the implied constant depends only on \mathbf{G} .

The natural sieving sets are

$$\Omega_\ell = \{g = (g_{\alpha,\beta}) \in \mathbf{G}(\mathbf{F}_\ell) \mid g_{i,j} \text{ is not a square in } \mathbf{F}_\ell\},$$

and by (2) of Proposition B.4 in Appendix B, we have

$$\frac{|\Omega_\ell|}{|G_\ell|} \gg 1$$

for $L \geq 3$ (for $L = 2$ and $\mathbf{G} = SL(2)$, the left-hand side vanishes), where the implied constant depends only on \mathbf{G} . (Note that the proof in Appendix B uses the Riemann Hypothesis over finite fields; the reader may find it interesting to see whether a more elementary argument may be found.)

Hence the sieve bound is

$$\mathbf{P}(\text{the } (i, j)\text{-th entry of } X_k \text{ is a square}) \ll (1 + L^{7d/4+1} \exp(-k\eta)) H^{-1}$$

with $H \gg L(\log L)^{-1}$ for $L \geq 3$, the implied constant depending on \mathbf{G} and S . As before, we take $L = \exp(4\eta/(7d + 4))$ if this is ≥ 3 and then obtain (7.16), and we deal with those k for which $\exp(4\eta/(7d + 4)) < 3$ by enlarging the implied constant. □

Remark 7.13 In the most classical sieves, estimating either the analogue of $R(\mathcal{L})$ or H is not a significant part of the work, the latter because once Ω_ℓ is known, which is usually not a problem there, it boils down to estimates for sums of multiplicative functions, which are well understood (see Appendix G). (See however the work of Duke [32] and Jones [68], where techniques of sums of Hurwitz class numbers and the trace formula are required to evaluate the size of the sifting sets.)

The results we have proved, and an examination of Appendix B, show that when performing a sieve in some group settings, sharp estimates for $R(\mathcal{L})$ or for H involve deeper tools. For the large sieve constant, this involves the representation theory of the group in non-trivial ways. For H , the issue of

estimating $|\Omega_\ell|$ may quickly become a difficult counting problem over finite fields. It is not hard to envision situations where the full force of Deligne's work on exponential sums over finite fields becomes really crucial, and not merely a convenience.

Remark 7.14 In the case of $SL(n, \mathbf{Z})$, the following trick of Rivin [108] shows that one can avoid the large sieve if one is interested only in a bound for the probability of having a reducible characteristic polynomial. Notice that if $g \in SL(n, \mathbf{Z})$ and $\det(T - g)$ is reducible, it has a non-trivial factor with constant coefficient ± 1 . So for any prime ℓ , the reduction of g is not in the set of matrices with a characteristic polynomial having a non-trivial factor of this type, which is easily checked to be of density $\ll \ell^{-1}$ (it is inside a bounded union of hypersurfaces), and choosing ℓ suitably after applying individual equidistribution (as in Remark 2.14) leads to a bound with exponential decay. However, note that this trick would not extend, say, to $SL(n, \mathbf{Z}_K)$ where \mathbf{Z}_K is the ring of integers in a number field K containing infinitely many units.

Exercise 7.2 Let $G = SL(2, \mathbf{Z})$ and let S be a finite symmetric generating set, (X_k) an associated random walk on G as in Proposition 7.8. Prove that for $k \geq 1$, we have

$$\mathbf{P}(X_k \text{ has a square coefficient}) \ll \exp(-ck^{1/3})$$

for some constant $c > 0$, c and the implied constant depending only on S .

Exercise 7.3 Say that a matrix $A \in M(n, \mathbf{R})$ is *strongly non-singular* if all its minors of all order $r \leq n$ are non-zero. Show that in the situation of Theorem 7.12, we have

$$\mathbf{P}(X_k \text{ is not strongly non-singular}) \ll \exp(-k\eta)$$

for some $\eta > 0$, where η and the implied constant depend only on \mathbf{G} and S .

From the first part of Theorem 7.12, we can easily deduce Theorem 1.3.

Corollary 7.15 Let $\mathbf{G} = SL(n)$, $n \geq 2$, or $Sp(2g)$, $g \geq 1$, let $G = \mathbf{G}(\mathbf{Z})$ and let $S = S^{-1}$ be a symmetric generating set of G , W the Weyl group of \mathbf{G} . Let (X_k) be a left-invariant random walk on G with independent uniformly distributed steps $\xi_k \in S$ such that

$$\mathbf{P}(\xi_k = s) = \mathbf{P}(\xi_k = s^{-1}) = p(s) > 0.$$

Then, almost surely, there exist only finitely many k such that $\det(T - X_k)$ is a polynomial with Galois group distinct from W , in particular such that

$\det(T - X_k)$ is reducible; or in other words, the set of matrices in G with reducible characteristic polynomial is transient for the random walk (X_k) .

Proof It suffices to apply the ‘easy’ Borel–Cantelli lemma (see Lemma F.2, (1), in Appendix F). Indeed, we have

$$\mathbf{P}(\det(T - X_k) \text{ has small Galois group}) \ll \exp(-\eta k)$$

for $k \geq 1$ by Theorem 7.12, if $n \geq 3$ (or $g \geq 2$), and therefore the series

$$\sum_{k \geq 0} \mathbf{P}(\det(T - X_k) \text{ has small Galois group})$$

obviously converges; if $n = 2$, the weaker bound in Exercise 7.2 is still more than enough to show that this still the case for $SL(2, \mathbf{Z})$. \square

Remark 7.16

- (1) Part of the point of this statement is that it *requires* some quantitative estimate for the probability that X_k has small Galois group (or reducible characteristic polynomial). Moreover, even if the distribution of the steps is uniform, it is not really possible to state this result coherently in the language of random products of some length N , since we wish to consider arbitrarily long walks, and the behaviour of the steps as we follow along this walk: this is a genuinely probabilistic statement.
- (2) As the second part of Exercise 6.3 shows, this transience phenomenon is also a reflection of the special properties of random walks on a non-commutative group such as $SL(n, \mathbf{Z})$.

7.6 Geometric applications

This section explains some applications of the ideas developed above to questions of geometry and topology. Such applications are quite appealing, and illustrate the potential relevance of some forms of sieve well outside analytic number theory.⁵

The first such result answers a question of Maher [96, Question 1.3], and was suggested by Rivin’s paper [108]. This has to do with the so-called mapping class groups of surfaces, and pseudo-Anosov elements in those groups, as defined by Thurston. We give a quick survey of the definitions involved in Section H.3 of

⁵ Other applications of sieve methods not directly related to arithmetic are already known, but they mostly involve unusual applications of *classical* sieves, e.g., the striking results of Goldfeld, Lubotzky and Pyber on counting congruence subgroups of arithmetic groups (see the description in [92, p. 120]), which involve the Bombieri–Vinogradov theorem on primes in arithmetic progressions to large moduli.

Appendix H. There is a nice and fairly detailed survey by Ivanov [65] and one of the standard references is the volume [40] by Fathi, Laudenbach and Poénaru, in particular Exposé 1 and 9 for the theory of pseudo-Anosov mapping classes.

Proposition 7.17 *Let G be the mapping class group of a compact connected orientable surface Σ_g of genus $g \geq 1$, let S be a finite symmetric generating set of G and let (X_k) , $k \geq 1$, be a left-invariant symmetric random walk on G with independent identically distributed steps (ξ_k) with*

$$\mathbf{P}(\xi_k = s) > 0 \text{ for all } s \in S.$$

Then the set $X \subset G$ of non-pseudo-Anosov elements is transient for this random walk.

Proof We follow the basic arguments of Rivin. The starting point is the existence of the surjective map

$$\rho : G \rightarrow Sp(2g, \mathbf{Z})$$

corresponding to the action of G on the first homology group $H_1(\Sigma_g, \mathbf{Z}) \simeq \mathbf{Z}^{2g}$ of the surface, preserving the non-degenerate alternating intersection pairing.

Let S be a generating set as above,⁶ and let $S' = \rho(S)$, a finite symmetric generating set for $Sp(2g, \mathbf{Z})$. The image $Y_k = \rho(X_k)$ of the random walk on G is a left-invariant random walk on $Sp(2g, \mathbf{Z})$ with steps $\xi'_k = \rho(\xi_k)$, distributed according to the image measure $\rho(\xi_k)$:

$$\mathbf{P}(\xi'_k = \rho(s)) = \sum_{\rho(t)=\rho(s)} \mathbf{P}(\xi_k = t) > 0,$$

for all $s \in S$. Hence (Y_k) is also symmetric, and Theorem 7.12 (if $g \geq 2$) or Proposition 7.8 (if $g = 1$, in which case ρ is in fact an isomorphism) applies to (Y_k) .

Now, the last (crucial) geometric point is the fact that it *suffices* that the following three conditions on the characteristic polynomial $P = \det(T - Y_k) = \det(T - \rho(X_k))$ hold for X_k to be pseudo-Anosov (this is the homological criterion for pseudo-Anosov diffeomorphisms, see [20, Lemma 5.1]):

- (i) P is irreducible;
- (ii) there is no root of unity which is a zero of P ;
- (iii) there is no $d \geq 2$ and polynomial Q such that $P(X) = Q(X^d)$.

⁶ It is a fairly deep fact that G is finitely generated, but observe that one can show more easily that there is a finitely generated subgroup mapping onto $Sp(2g, \mathbf{Z})$, and one could argue for any such subgroup instead...

Accordingly we have

$$\mathbf{P}(X_k \text{ is not pseudo-Anosov}) \leq p_1 + p_2 + p_3$$

where p_1, p_2, p_3 are the probabilities that $\det(T - Y_k)$ satisfies those three conditions.

Assume first $g \geq 2$. Then, by Part (1) of Theorem 7.12, there exists $\eta_1 > 0$, depending only on S and g such that

$$p_1 \ll \exp(-\eta_1 k)$$

for $k \geq 1$. To estimate p_2 and p_3 , we can use simpler sieves (or even merely individual equidistribution, because the sifting sets will have density going to zero with ℓ) to obtain comparable bounds. For p_2 , since P is an integral polynomial of degree $2g$ and hence may only have roots of unity with bounded order as zeros, it suffices to estimate the probability of the sifted set associated to the sieving sets

$$\Omega_\ell = \{g \in Sp(2g, \mathbf{F}_\ell) \mid (\Phi_d \pmod{\ell}) \nmid \det(T - g)\}$$

for some fixed $d \geq 1$, where $\Phi_d \in \mathbf{Z}[X]$ is the d -th cyclotomic polynomial. We have $|\Omega_\ell| \gg |Sp(2g, \mathbf{F}_\ell)|$, in fact $|\Omega_\ell| \sim |Sp(2g, \mathbf{F}_\ell)|$ (see Lemma B.5 in Appendix B), hence the sieve again yields $p_2 \ll \exp(-\eta_2 k)$ for $k \geq 1$ and some constant $\eta_2 > 0$ (depending only on g and S).

For p_3 , we consider similarly

$$\Omega'_\ell = \{g \in Sp(2g, \mathbf{F}_\ell) \mid \det(T - g) \text{ is not of the form } Q(X^d)\}$$

for some fixed $d \geq 2$. We also have $|\Omega'_\ell| \gg |Sp(2g, \mathbf{F}_\ell)|$ rather trivially, and $p_3 \ll \exp(-\eta_3 k)$ for some constant $\eta_3 > 0$.

Now we conclude that

$$\mathbf{P}(X_k \text{ is not pseudo-Anosov}) \ll \exp(-\eta k)$$

for $\eta = \min(\eta_1, \eta_2, \eta_3)$, and we can again apply the Borel–Cantelli lemma as in the proof of Corollary 7.15.

The argument with $g = 1$ is exactly similar, appealing to Proposition 7.8 which gives weaker bounds, more than sufficient to obtain the desired transience. \square

Remark 7.18 Maher [96] proved that the probability that X_k is pseudo-Anosov tends to 1 as $k \rightarrow +\infty$ using rather more information concerning the geometry and structure of the mapping class group (in particular, more about pseudo-Anosov classes, beyond the ‘negative’ homological criterion), and the limiting behaviour of the random walks. His methods did not lead

to a quantitative bound for the probability that (X_k) is not pseudo-Anosov, hence didn't answer the question of transience. However, it is important to note that his result is also more general, and applies to random walks on any subgroup of G which is not 'obviously too small' in some sense. It should be emphasized that this condition encompasses groups which seem utterly out of reach of the sieve, for instance the Torelli group T_g which is the kernel of the homology action ρ . It may seem surprising (for beginners, such as the author . . .) that pseudo-Anosov mapping classes actually exist in this subgroup, but Maher's result shows that they remain 'generic' (see [40, p. 250] for a construction which gives some examples, and the observation that Nielsen had conjectured they did not exist). It would be interesting to know (using sieve or otherwise) if a random walk on the Torelli group is still transient on the set of pseudo-Anosov elements, or if there is a genuine difference in behaviour of this subgroup.

Exercise 7.4 Maher [96] gives some further examples of properties of elements of mapping class groups which have probability going to 1 as the length of a random walk goes to infinity. This exercise sketches a proof of a 'transience' form of a property which is slightly weaker than one he considers.

- (1) Let $g \geq 1$ and $T \geq 1$ be fixed. Show that there are only finitely many irreducible monic polynomials $P \in \mathbf{Z}[X]$ of degree $2g$ such that $P(0) = 1$ and such that all roots ρ of P in \mathbf{C} satisfy $|\rho| \leq T$. (This number of course depends on g and T .)
- (2) Let $A \subset Sp(2g, \mathbf{Z})$ be the set of those matrices g such that $\det(T - g)$ satisfies the three conditions (i), (ii), (iii) in the proof of Proposition 7.17. Deduce from (1) that the set

$$A_T = \{g \in A \mid \text{all roots of } \det(T - g) \text{ are of modulus } \leq T\}$$

is transient for a symmetric left-invariant random walk on $Sp(2g, \mathbf{Z})$ with respect to a finite symmetric generating set as before. [Hint: In addition to the previous sieve, sieve by excluding those $g \in Sp(2g, \mathbf{F}_\ell)$ with characteristic polynomial equal to the reduction of one of the finitely many polynomials of the previous question.]

- (3) Deduce that (for fixed $T \geq 1$), in a symmetric random walk (X_k) on the mapping class group G of a closed surface S of genus $g \geq 1$ with respect to a symmetric generating set, the set of elements f which are either not pseudo-Anosov, or pseudo-Anosov with dilation factor $\lambda(f) < T$, is transient. In particular, the dilation factor (or expanding factor) $\lambda(X_k)$ tends to infinity almost surely in such a random walk (where λ is extended to $f \in G$

which are not pseudo-Anosov by setting $\lambda(f) = 0$ for those mapping classes). [Hint: See [40, Exposé 9] for the definition of the dilation factor⁷ of a pseudo-Anosov class $f \in G$, and in particular [40, Theorem, p. 190; Proposition, p. 194] for the lower bound $\lambda(f) \geq \gamma_{\rho(f)}$, where γ_g is the largest modulus of a root of $\det(T - g)$.]

- (4) Prove a quantitative bound for the rate of divergence of $\lambda(f)$. (Note that such a bound from the above proof may be far from the truth because we are detecting an ‘Archimedean’ condition, namely that some real number is less than T , by means of reduction modulo primes . . .)

Exercise 7.5 Rivin [108, Section 10] gives another application of sieve ideas, which is similar in spirit, to the ‘generic’ behaviour of automorphisms of free groups. Let F_n be the free group on $n \geq 2$ generators x_1, \dots, x_n , and let $G = \text{Aut}(F_n)$. This is a finitely generated group (in fact, finitely presented), as proved by Nielsen, and indeed the automorphisms

$$\alpha_i : \begin{cases} x_i \mapsto x_i^{-1} \\ x_j \mapsto x_j, \text{ for } j \neq i, \end{cases} \quad \beta_{i,j}^{\pm} : \begin{cases} x_i \mapsto x_i x_j^{\pm 1} \\ x_k \mapsto x_k, \text{ if } k \neq i \end{cases}$$

for $1 \leq i \leq n$, $j \neq i$, form a symmetric generating set S (see, e.g., [95, Proposition 4.1]). The elements of G act on the abelianization $G/[G, G] = \mathbf{Z}^n$, giving a map

$$\rho : G \rightarrow GL(n, \mathbf{Z})$$

which is onto (this is clear from the fact that S maps to a generating set of $GL(n, \mathbf{Z})$, indeed $\rho(\alpha_i)$ is the reflection with axis the i -th coordinate axis, and $\rho(\beta_{i,j})$ is an elementary matrix $E_{i,j}(1)$).

- (1) By setting up a coset sieve⁸ corresponding to the exact sequence

$$1 \rightarrow SL(n, \mathbf{Z}) \rightarrow GL(n, \mathbf{Z}) \rightarrow \{\pm 1\} \rightarrow 1,$$

show that in a symmetric random walk (X_k) on G with respect to a finite symmetric set of generators, the set N of automorphisms $\alpha \in G$ such that $\det(T - \rho(\alpha^k))$ is reducible for some $k \geq 1$ is transient. [Hint: Show that this stronger form of irreducibility is implied by the Galois group of the splitting field of the characteristic polynomial $\det(T - \rho(\alpha))$ being the full symmetric group, then argue as in Gallagher’s Theorem 4.2.]

⁷ ‘Rapport de dilatation’ in French.

⁸ This is done in the generality of this chapter in F. Jouve’s thesis, [69].

- (2) Deduce that in a random walk (X_j) as above, the set Y of automorphisms which do not have the *strong irreducibility*, or *iwip* property,⁹ is transient, where an element $\alpha \in G$ is strongly irreducible if and only if there is no $k \geq 1$ such that α^k sends a free factor H of F_n to a conjugate of itself.

Another fairly direct geometric application of the large sieve for random walks in mapping class groups arises from work of N. Dunfield and W. Thurston. In their paper [34], they define a notion of ‘random 3-manifold’ and study some of its properties with respect (among other things) to the existence of finite Galois coverings with certain Galois groups, especially with regard to homological properties, such as having positive first Betti number.

Again let $g \geq 1$ be an integer, and let G denote the mapping class group of a closed orientable surface Σ_g of genus g , with a fixed finite symmetric set of generators S ; for $g = 1$ (in which case $G = SL(2, \mathbf{Z})$), assume that $1 \in S$ to avoid the periodicity issues. Then let (X_k) for $k \geq 0$ denote a random walk on G given by independent symmetric increments ξ_k .

Associated to this random walk, Dunfield and Thurston define a sequence (M_k) of random 3-manifolds by the following process, known as ‘Heegaard splitting’: M_k is obtained from two copies of a handlebody¹⁰ H_g of genus g with boundary $\partial H_g = \Sigma_g$ by identifying their common boundary using a diffeomorphism in the mapping class $X_k \in G$; one shows that this manifold, up to diffeomorphism, depends only on the class X_k . It is a fact from topology¹¹ that *any* compact, orientable, connected 3-manifold can be obtained by such a process (non-uniquely), for *some* genus $g \geq 1$. However, although this goes a long way towards justifying the relevance of the random manifolds M_k if one wishes to know something about what to expect from general 3-manifolds, it is not necessary for what follows.

Dunfield and Thurston¹² study properties of the fundamental group $\pi_1(M_k, x_0)$ of M_k (with respect to an arbitrary base-point), motivated by the so-called Virtual Haken Conjecture, which states (geometrically) that any orientable compact connected 3-manifold M with infinite fundamental group has a finite covering $N \rightarrow M$ such that the first Betti number $b_1(N) = \dim_{\mathbf{Q}} H_1(N, \mathbf{Q})$ of N is > 0 , or equivalently (algebraically) that there exists

⁹ ‘Irreducible With Irreducible Powers’.

¹⁰ That is, a ‘filled’ doughnut with g holes; see Appendix H.

¹¹ Actually, a very old one: Heegaard introduced this idea in his 1898 dissertation.

¹² For the number theorist, there is a distinct flavour of Cohen–Lenstra heuristics in their paper.

a finite index subgroup $H \subset \pi_1(M)$ with infinite abelianization.¹³ This conjecture seems to be the most important open question concerning 3-manifolds (now that the Poincaré and the geometrization conjecture are considered to be proved).

In particular, the following questions are then very natural for the random 3-manifolds described above:

- What is the probability that $\pi_1(M_k)$ has a finite index normal subgroup H with $\pi_1(M_k)/H$ isomorphic to a given finite group Q ?
- If such a finite index subgroup exists, corresponding to a finite covering $N \rightarrow M_k$, what is the probability that the first Betti number of N is positive?

As mentioned by Dunfield and Thurston, one could hope to find this probability to be positive (for k and/or g large enough) to provide many instances of manifolds for which the Virtual Haken Conjecture holds. However, Dunfield and Thurston provide strong evidence that this probability is asymptotically zero when $k \rightarrow +\infty$. If true, the Virtual Haken Conjecture seems to be quite hard to catch.

As in Section 8 of [34], we will look at the structure of the first homology group $H_1(M_k, \mathbf{Z})$ of the manifolds M_k . We recall (again, as done in Appendix H) the following properties of the first homology group of an arbitrary orientable, compact, connected manifold M (of dimension not necessarily equal to 3):

- This group is an abelian group of finite type, and is in fact the abelianization of the fundamental group $\pi_1(M)$, and as such, it classifies the coverings of M_k with abelian Galois group (this shows that the problem is related to the considerations of random groups in Section 6.2).
- The first homology group with rational coefficients, $H_1(M, \mathbf{Q})$, is isomorphic to $H_1(M, \mathbf{Z}) \otimes \mathbf{Q}$ and is therefore a finite dimensional \mathbf{Q} -vector space, of dimension (namely, the first Betti number of M) at most equal to the rank of $H_1(M, \mathbf{Z})$.
- For any prime ℓ , the first homology group $H_1(M, \mathbf{F}_\ell)$ with coefficients in the finite field \mathbf{F}_ℓ is isomorphic to $H_1(M, \mathbf{Z}) \otimes \mathbf{F}_\ell = H_1(M, \mathbf{Z})/\ell H_1(M, \mathbf{Z})$ and is therefore a finite dimensional \mathbf{F}_ℓ -vector space, of dimension at least equal to the rank of $H_1(M, \mathbf{Z})$.

With this setup, the results are as follows: Dunfield and Thurston show (see Corollary 8.5 in [34]) that, given a prime number ℓ , the probability that the group $H_1(M_k, \mathbf{F}_\ell)$ is zero tends to 1 as $k \rightarrow +\infty$ and $\ell \rightarrow +\infty$, and in particular,

¹³ This equivalence is a consequence of the link between coverings of M and its fundamental group, and of the fact that the first Betti number is the rank as abelian group of the abelianization of the fundamental group; see Appendix H.

the probability that $H_1(M_k, \mathbf{Q}) \neq 0$ tends to 0, since the first Betti number is at most the dimension of $H_1(M_k, \mathbf{F}_\ell)$. Contrariwise, they show that the *expected value* of the order of $H_1(M_k, \mathbf{Z})$ tends to infinity as $k \rightarrow +\infty$ (this group is finite whenever $H_1(M_k, \mathbf{Q}) = 0$). These results are of course comparable with Proposition 6.4 (also derived in qualitative form in [34]).

The same arguments together with the sieve easily yield the following quantitative results:

Proposition 7.19 *Let $g \geq 1$ be given and let (M_k) be a sequence of random 3-manifolds as defined above. Then*

(1) *There exists $\delta > 0$ such that*

$$\mathbf{P}(H_1(M_k, \mathbf{Q}) \neq 0) \ll \exp(-\delta k) \quad (7.20)$$

for all $k \geq 1$, the implied constant as well as δ depending only on g , S and the distribution of the steps ξ_k of the underlying random walk. In particular, the set of all 3-manifolds with positive first Betti number is transient.

(2) *There exist $b > 0$, $\alpha > 0$, $C \geq 0$ and $C' \geq 0$ such that*

$$\mathbf{P}\left(H_1(M_k, \mathbf{F}_\ell) \neq 0 \text{ for at least } \log bk \text{ primes}\right) \geq 1 - \frac{C}{\log k}, \quad (7.21)$$

$$\mathbf{P}\left(\text{The order of } H_1(M_k, \mathbf{Z}) \text{ is } < k^{\alpha \log \log k}\right) \leq \frac{C'}{\log k}, \quad (7.22)$$

and in particular we have

$$\mathbf{E}\left(\text{Order of } H_1(M_k, \mathbf{Z})_{\text{tors}}\right) \gg k^{\alpha \log \log k}$$

where $H_1(M_k, \mathbf{Z})_{\text{tors}}$ is the torsion subgroup of $H_1(M_k, \mathbf{Z})$. The constants b , α , C and C' as well as the implied constant depend only on g , S and the distribution of ξ_k .

This shows that with probability going to 1, $H_1(M_k, \mathbf{Z})$ is a finite abelian group with ‘superpolynomial’ growth in terms of k . Since (7.22) will be deduced rather wastefully from (7.21), it is even possible that the size of $H_1(M_k, \mathbf{Z})$ could be exponentially growing. On the other hand, it’s not clear how to trade a faster convergence of the probability in (2) for a slower growth of $H_1(M_k, \mathbf{Z})$.

Remark 7.20 The comparison with Section 6.2 is instructive, and bears on the important philosophical question which asks what special properties distinguish

fundamental groups of 3-manifolds from general finitely presented groups.¹⁴ Dunfield and Thurston find strong evidence that the fundamental groups of M_k seem to have more finite quotients than the random finitely presented groups described in Section 6.2 (in terms of asymptotic probabilities for the existence of a quotient isomorphic to a given group Q , for instance). Our results are contrasted in this respect: on the one hand, the lower bound in part (1) of Proposition 6.4 together with (1) of Proposition 7.19 shows that the random groups have much higher (though small) probability of having infinite abelianization – polynomial decay instead of exponential decay. But on the other hand, when the abelianization is finite, it tends to be much larger for 3-manifolds than for random groups (for the latter, we observed that the trivial bound for the expected size of the determinant of the matrix giving the size of the abelianization is $\ll k^g$, compared with the superpolynomial growth of Proposition 7.19).

It is not clear to the author what should be thought of this. Perhaps only asymptotic properties of the probabilities are relevant to the comparison? Or perhaps, in fact, one should use more sophisticated types of random groups to compare with the 3-manifolds? For instance, a notion of random groups introduced by Gromov has the property of yielding groups which have Property (T) with probability converging to 1 exponentially fast, as proved by Silberman (see [121, Theorem 2.16]). Hence these groups have finite abelianization with at least this probability (a basic property of groups with Property (T), see Appendix D). However, the construction of these groups is rather more sophisticated than the one in Section 6.2: they are quotients of free groups where the relations are obtained from words arising by following a random labelling of edges of a graph, which is ultimately taken from a family of expanders. In particular, the number of relations is not fixed but grows with the size of the group, which indicates that this model is not a good comparison point for fundamental groups of 3-manifolds. (Which, in any case, can probably not have Property (T) unless they are finite, as this would contradict the Virtual Haken Conjecture; in fact, Lubotzky and Sarnak conjecture that fundamental groups of hyperbolic 3-manifolds do not even have Property (τ).)

The proof of Proposition 7.19 proceeds by combining the analysis of $H_1(M_k, \mathbf{Z})$ in [34] with applications of equidistribution and of the large sieve

¹⁴ This is an issue only in dimension 3; in dimensions 1 and 2, fundamental groups of compact manifolds are entirely classified and rather well understood, while in dimension $d \geq 4$, it is known that *any* finitely presented group arises as the fundamental group of an orientable compact connected manifold of dimension d without boundary, see, e.g. [57, V.27–V.29]. In particular, there are examples of manifolds with infinite fundamental group where no finite cover has positive first Betti number – it suffices to take a manifold with fundamental group infinite and having Property (T).

(as was the case for Proposition 6.4) with the sieve setting given by the group sieve

$$(Sp(2g, \mathbf{Z}), \{\text{primes}\}, Sp(2g, \mathbf{Z}) \rightarrow Sp(2g, \mathbf{F}_\ell))$$

with the siftable set $(\Omega, \mathbf{P}, \rho(X_k))$, where ρ is, as before, the map giving the homology action of a mapping class. Indeed, we have the following lemma found in [34, Section 8] which describes the groups $H_1(M_k, \mathbf{Z})$ and $H_1(M_k, \mathbf{F}_\ell)$ in terms of the given surface Σ_g and the homology action of the mapping class X_k defining M_k :

Lemma 7.21 *Let $\varphi \in G$ be a mapping class and let M_φ be the 3-manifold obtained by gluing two copies of H_g along their common boundary Σ_g using the mapping class φ .*

(1) *Let $J = \text{Ker}(H_1(\Sigma_g, \mathbf{Z}) \rightarrow H_1(H_g, \mathbf{Z}))$. Then*

$$H_1(M_\varphi, \mathbf{Z}) \simeq H_1(\Sigma_g, \mathbf{Z}) / \langle J, \rho(\varphi)^{-1}(J) \rangle,$$

and moreover $J \simeq \mathbf{Z}^g$ is a Lagrangian sublattice in $H_1(\Sigma_g, \mathbf{Z})$ with respect to the intersection pairing; in other words, J is a lattice of rank g , and the intersection pairing is identically zero when restricted to J .

(2) *For any prime ℓ , we have similarly*

$$H_1(M_\varphi, \mathbf{F}_\ell) \simeq H_1(\Sigma_g, \mathbf{F}_\ell) / \langle J_\ell, \rho_\ell(\varphi)^{-1}(J_\ell) \rangle$$

where $J_\ell = J/\ell J$ is the image of J in $H_1(\Sigma_g, \mathbf{F}_\ell) \simeq H_1(\Sigma_g, \mathbf{Z}) \otimes \mathbf{F}_\ell \simeq \mathbf{F}_\ell^{2g}$.

Proof of Proposition 7.19 Since the handlebody H_g and the boundary surface Σ_g are fixed throughout the argument, the lattice $V = H_1(\Sigma_g, \mathbf{Z})$ and its Lagrangian sublattice J (given by the lemma) are likewise fixed, as well as their reductions $J_\ell \subset V_\ell$ modulo ℓ . Now let

$$\Omega_\ell = \{x \in Sp(V_\ell) \mid \langle J_\ell, x^{-1}(J_\ell) \rangle \neq \mathbf{F}_\ell^{2g}\} \subset Sp(V_\ell) \simeq Sp(2g, \mathbf{F}_\ell)$$

be the set of symplectic matrices over \mathbf{F}_ℓ for which the two Lagrangian subspaces J_ℓ and $x^{-1}(J_\ell)$ are not transverse.

By the lemma applied to $\varphi = X_k$, we have the basic criterion

$$H_1(M_k, \mathbf{F}_\ell) \neq 0 \text{ if and only if } \rho_\ell(X_k) \in \Omega_\ell \quad (7.23)$$

which allows us to reduce the statements of Proposition 7.19 to sieve conditions.

We start with part (1). As recalled above, we have the basic upper bound

$$\dim_{\mathbf{Q}} H_1(M_k, \mathbf{Q}) \leq \text{rank } H_1(M_k, \mathbf{Z}) \leq \dim_{\mathbf{F}_\ell} H_1(M_k, \mathbf{F}_\ell),$$

and hence, if $\dim H_1(M_k, \mathbf{Q}) \geq 1$, it follows from the criterion that $\rho_\ell(X_k) \in \Omega_\ell$ for any prime ℓ .

We use equidistribution for the fixed quotient $Sp(2g, \mathbf{F}_\ell)$ (see Remark 2.14): using Proposition 7.2 (and Lemma 7.3), we find easily that for any prime ℓ , we have

$$\mathbf{P}(H_1(M_k, \mathbf{Q}) \neq 0) \leq \mathbf{P}(\rho_\ell(X_k) \in \Omega_\ell) = \frac{|\Omega_\ell|}{|Sp(2g, \mathbf{F}_\ell)|} + O(\ell^A \exp(-k\eta))$$

for some constant $A \geq 0$ and $\eta > 0$, which depend only on g , S and the distribution of the steps ξ_k of the random walk.

The density above is computed exactly in [34, Section 8.3]; namely we have

$$\frac{|\Omega_\ell|}{|Sp(2g, \mathbf{F}_\ell)|} = 1 - \prod_{j=1}^g \frac{1}{1 + \ell^{-j}} \quad (7.24)$$

(for completeness we sketch the proof in Proposition B.4, (7), of Appendix B).

From this, using a Taylor expansion at 0 of

$$x \mapsto \frac{1}{1+x} \cdots \frac{1}{1+x^g}.$$

we find that

$$\frac{|\Omega_\ell|}{|Sp(2g, \mathbf{F}_\ell)|} = \frac{1}{\ell} + O\left(\frac{1}{\ell^2}\right)$$

for $\ell \geq 2$ (and fixed g).

Taking $\ell \asymp \exp(k\eta/(A+1))$ (if this is ≥ 2 ; otherwise, the bound (7.20) is trivial anyway, by increasing the constant C if need be), we obtain

$$\mathbf{P}(H_1(M_k, \mathbf{Q}) \neq 0) \ll \frac{1}{\ell} \ll \exp\left(-\frac{k\eta}{A+1}\right),$$

the implied constant depending on g , S and the distribution of the steps of the random walk.

To deal with part (2), we use the dual sieve (Proposition 2.15, with \mathcal{L}^* the set of primes $\leq L$) as in Proposition 6.4. Tautologically, the criterion (7.23) yields

$$H_1(M_k, \mathbf{F}_\ell) = 0 \text{ if and only if } \rho_\ell(X_k) \notin \Omega_\ell. \quad (7.25)$$

We obtain then by (2.13) the estimate

$$\mathbf{E}\left((P(X_k, L) - P(L))^2\right) \leq (1 + L^A \exp(-k\eta))P(L)$$

(where $A \geq 0$ and $\eta > 0$ are not necessarily the same as before, but again depend only on g , S and the distribution of the steps of the random walk), and

$$P(X_k, L) = \sum_{\substack{\ell \leq L \\ X_k \pmod{\ell} \in \Omega_\ell}} 1, \quad P(L) = \sum_{\ell \leq L} \frac{|\Omega_\ell|}{|Sp(2g, \mathbf{F}_\ell)|}.$$

Notice that by the properties of the first homology groups with coefficients in finite fields and by (7.25), $P(X_k, L)$ is equal to the number of primes $\ell \leq L$ for which $H_1(M_k, \mathbf{F}_\ell) \neq 0$. So this inequality means that if L is small enough, this number will be close to $P(L)$ with high probability.

The density bound above for Ω_ℓ gives

$$P(L) \gg \log \log L$$

for $L \geq 4$. By positivity, we write

$$\mathbf{E}\left((P(X_k, L) - P(L))^2\right) \geq \frac{1}{4} P(L)^2 \mathbf{P}\left(P(X_k, L) < \frac{1}{2} P(L)\right).$$

Let L_0 be large enough that we have $P(L) \geq (\log \log L)/2$ for all $L \geq L_0$ (L_0 exists, and depends only on g). Then for $L \geq L_0$, we obtain

$$\begin{aligned} \mathbf{P}(P(X_k, L) < \tfrac{1}{4} \log \log L) &\leq (1 + L^A \exp(-k\eta)) P(L)^{-1} \\ &\leq 2 \frac{1 + L^A \exp(-k\eta)}{\log \log L}. \end{aligned}$$

We select $L = \exp(k\eta/A)$, if this is $\geq L_0$ (otherwise the estimate (7.21) is trivial after increasing the constant C), and obtain that

$$\mathbf{P}(P(X_k, L) < \tfrac{1}{4} \log bk) \leq \frac{4}{\log bk}.$$

with $b = \eta/A$. In other words, with probability at least $1 - 4(\log bk)^{-1}$, $H_1(M_k, \mathbf{F}_\ell) \neq 0$ for at least $(\log bk)/4$ distinct primes, which implies (7.21).

Now to go from this to the lower bound (7.22) for the size of $H_1(M_k, \mathbf{Z})$, we argue as follows: if $H_1(M_k, \mathbf{Z})$ is finite, and if $P(X_k, L) \geq (\log bk)/4$, then $H_1(M_k, \mathbf{Z})$ has non-zero reduction modulo ℓ for at least that many primes, and its size is at least the product of those primes. We don't know how the primes which occur are distributed, but the product involved is at least as large as the product of the first $[(\log bk)/4]$ primes. Thus with probability at least $1 - 4/(\log bk)$, we have

$$|H_1(M_k, \mathbf{Z})| \geq \prod_{\ell \leq U} \ell$$

where U is the $[(\log bk)/4]$ -th prime. Using easy Chebychev estimates, the k -th prime is $\gg k \log k$ (for $k \geq 2$) and the sum of logarithms of primes $\leq U$ is $\gg U$

for $U \geq 2$, so we have first $U \gg (\log bk)(\log \log bk)$, and

$$\prod_{\ell \leq U} \ell = \exp \left(\sum_{\ell \leq U} \log \ell \right) \geq \exp(f(\log bk)(\log \log bk)) \geq k^{\alpha \log \log k}$$

for some $f > 0$ and $\alpha > 0$. Therefore, we have shown that

$$\mathbf{P}(\text{Order of } H_1(M_k, \mathbf{Z}) < k^{\alpha \log \log k}) \leq \frac{4}{\log bk},$$

hence (7.22) follows. \square

Note that part (2) may again be compared with the fact that ‘almost all’ integers $n \leq x$ have about $\log \log x$ prime divisors (counted without multiplicity). Because it is easy to see that the matrices $\rho(X_k)$ have coefficients of size at most exponential in k , it follows straightforwardly from Lemma 7.21 that the size of the torsion group of $H_1(M_k, \mathbf{Z})$ is also at most of exponential size, and therefore the logarithmic order of magnitude of the number of prime factors we found (with high probability) is best possible.

If the order of $H_1(M_k, \mathbf{Z})$ (or of its torsion part rather) behaves like a ‘random’ integer, we would expect that the presence of roughly $\log k$ prime divisors (with large probability) implies that this integer is indeed of size exponential in k . However, the author lacks geometric and topological experience to have any idea if this ‘randomness’ is a reasonable expectation.

In another paper, Dunfield and Thurston [35] present experimental evidence coming from a database which contains 10 986 distinct hyperbolic 3-manifolds. Looking at this data set, the maximal size of the torsion subgroup of $H_1(M, \mathbf{Z})$ is 423, and the histogram of the values of the size of the first homology group looks roughly like that of an exponential distribution (with mean approximately 62.92791); however, the number of prime factors doesn’t exceed five, and hence it’s unclear how meaningful a comparison between the experimental data and the number of prime factors of integers sampled according to an approximation to this exponential distribution can be.

Another point is that Dunfield and Thurston observe [34, 9.1] that a large majority (roughly 8000 of them) of the manifolds in their census have a fundamental group with two generators, and hence can probably be obtained by Heegaard splitting with $g = 2$. This is interesting in terms of the comparison between random finitely presented groups and fundamental groups, since by the last remark in (1) of Proposition 6.4, the set of 2-generator groups with two relations with infinite abelianization was found to be *recurrent*.

Problem 7.22 Maher has shown [96], using work of Hempel and the geometrization conjecture, that M_k is hyperbolic with probability tending to 1 as k tends to infinity. Can you prove that the set of non-hyperbolic manifolds is transient for a random 3-manifold as above?

7.7 Explicit bounds and arithmetic transitions

Classically, an important feature of sieve methods has been their uniformity and the explicitness of the results. In the applications of this chapter, this aspect is somewhat diminished in general because the evaluation of the large sieve constant involves the Kazhdan or Property (τ) constants of the discrete group, which are simply asserted to exist. This illustrates that it would be highly interesting to know such constants explicitly. As mentioned previously, this is a question in harmonic analysis or geometric group theory which was first raised by Serre, de la Harpe and Valette.

The first such results were proved by Burger, who for $G = SL(3, \mathbf{Z})$ gave (among other things) explicit (τ)-constants for representations factoring through a quotient $SL(3, \mathbf{Z}/m\mathbf{Z})$ (see [18] and the Appendix in [58]). Then an important breakthrough was the work of Shalom [118], who (among other things!) found explicit Kazhdan constants for $SL(n, \mathbf{Z})$ with respect to the symmetric generating set S of elementary matrices $E_{i,j}(\pm 1)$ with ± 1 in the (i, j) -th entry; see Theorem D.1 in Appendix D for a sketch of this result. Building on this work, Kassabov recently obtained even stronger bounds which are, in a sense, best possible. Using this we can obtain explicit sieve bounds for random walks on $SL(n, \mathbf{Z})$ with respect to this generating set, and even keep control of uniformity with respect to n .

Proposition 7.23 *Let $n \geq 3$ be an integer, let S be the generating set of $G = SL(n, \mathbf{Z})$ defined above, and let (X_k) be the symmetric left-invariant random walk on G with independent steps ξ_k uniformly distributed according to*

$$\mathbf{P}(\xi_k = s) = \frac{1}{|S|} = \frac{1}{2(n^2 - n)}$$

for all $s \in S$.

(1) *The estimates (7.11) and (7.12) hold with*

$$\begin{aligned} \eta &= \eta_n = -\log\left(1 - \frac{1}{8n(n-1)(21\sqrt{n} + 460)^2}\right) \\ &\geq \frac{1}{8n(n-1)(21\sqrt{n} + 460)^2}. \end{aligned}$$

(2) We have

$$\mathbf{P}(\det(T - X_k) \text{ has small splitting field}) \ll k \exp(-k \frac{\eta_n}{n^2})$$

for all $k \geq 1$ and all $n \geq 3$, the implied constant being absolute. Recall that an integral polynomial of degree n has small splitting field if the Galois group of the splitting field over \mathbf{Q} is not \mathfrak{S}_n .

We give a precise value of η simply for pleasure. Note that $\eta \sim (3528n)^{-3}$ as $n \rightarrow +\infty$ and that

$$\eta \geq \frac{1}{14112n^2(n-1)}, \quad \text{if } n \geq 480.$$

Proof

- (1) According to the proof of Theorem 7.4, and since we have $p^+ = \min \mathbf{P}(\xi_k = s) = 1/|S|$, the bounds we seek are valid with

$$\eta = \eta_n = \min \left(-\log \frac{1}{1 - \frac{\kappa}{4(n^2-n)}}, -\log \frac{1}{1 - \frac{1}{(n^2-n)c^2}} \right)$$

as given by (7.5) (since $p^+ = 1/|S|$ here), where c is the length of an S -relation of odd length in G , and κ is the Kazhdan constant for S .

First of all, the commutator relation

$$E_{1,2}(1)[E_{1,3}(1), E_{3,2}(1)]^{-1} = 1$$

(which uses $n \geq 3$, by the way) implies that we can take $c = 5$ (and it is easy to see that this is the best result; there is no loop of length 3).

Much more deeply, Kassabov's result [70] states that for any unitary representation π of G not containing the trivial representation, and any non-zero vector v in the space of π , there exists $s \in S$ such that $\|\pi(s)v - v\| \geq \varepsilon_n \|v\|$, with $\varepsilon_n = (42\sqrt{n} + 920)^{-1}$. This means we can take $\kappa = \varepsilon_n^2$, which is of size roughly $1/(1764n)$. It is clear that the smallest in the two quantities defining η is the one involving κ , hence the result.

- (2) In order to derive a result which is uniform with respect to n , we repeat the basic steps of the proof of Theorem 7.12, as in the proof of the uniform version of Gallagher's Theorem. What is needed is a uniform estimate for $R(\mathcal{L})$ instead of (7.18), the 'right' choice of conjugacy classes to distinguish the symmetric group from its subgroups, and a uniform lower bound for the corresponding sums H instead of (7.19).

We start with the first point; taking $\mathcal{L} = \{3 \leq \ell \leq L-1\}$ for convenience, Lemma 5.9 gives

$$A_1(SL(n, \mathbf{F}_\ell)) \leq n(\ell + 1)^{(n^2-n)/2+n} \frac{\ell + 1}{\ell - 1} \leq 2nL^{(n^2+n)/2},$$

$$A_\infty(SL(n, \mathbf{F}_\ell)) \leq (\ell + 1)^{(n^2-n)/2} \leq L^{(n^2-n)/2}$$

and therefore

$$R(\mathcal{L}) = \max_{3 \leq \ell \leq L-1} A_\infty(SL(n, \mathbf{F}_\ell)) \sum_{3 \leq \ell \leq L-1} A_1(SL(n, \mathbf{F}_\ell)) \leq 2nL^{n^2-n+1}.$$

We next combine the sieves using the sets $\Omega_{i,\ell} \subset SL(n, \mathbf{F}_\ell)$ of matrices with characteristic polynomial f which:

- is irreducible for $\Omega_{1,\ell}$;
- is a product of an irreducible quadratic polynomial and other distinct irreducible polynomials of odd degree for $\Omega_{2,\ell}$;
- has an irreducible factor of prime degree $p > n/2$ for $\Omega_{3,\ell}$.

By the argument of Bauer, the probability that $\det(T - X_k)$ has small splitting field is at most the sum of the probabilities of the corresponding sifted sets.

By combining Gallagher's bounds for the density δ_i of the conjugacy classes in \mathfrak{S}_n corresponding to the above splitting types (see (4.3) and (4.4)), Lemma B.2 which gives a precise lower bound for the number of monic polynomials of degree n with constant term 1 of each splitting type, and Proposition B.4, (1), we deduce that for $i = 1, 2, 3$ we have

$$\begin{aligned} \frac{|\Omega_{i,\ell}|}{|SL(n, \mathbf{F}_\ell)|} &\geq \delta_i \left(1 - \frac{1}{\ell}\right)^{n^2+1} \left(1 - \frac{1}{\ell}\right)^{2n} \left(1 - \frac{1}{\sqrt{\ell}}\right)^n \\ &\geq \delta_i \left(1 - \frac{1}{\sqrt{\ell}}\right)^{4n^2} \end{aligned}$$

for $\ell > 16n^2$. Therefore, if $L > 16n^2$, we find

$$H \geq \delta_i \sum_{16n^2 < \ell \leq L-1} \left(1 - \frac{1}{\sqrt{\ell}}\right)^{4n^2}.$$

By the mean-value theorem we have

$$\left(1 - \frac{1}{\sqrt{\ell}}\right)^{4n^2} = 1 + O\left(\frac{n^2}{\sqrt{\ell}}\right)$$

for all $\ell \geq 3$, with an absolute implied constant (in fact, it is at most 1), and since δ_i is smallest for $i = 1$, where it is $1/n$, it follows that

$$H \gg \frac{\pi(L-1)}{n} + O(n^2 \sqrt{L} (\log L)^{-1})$$

for $L > 16n^2 \geq 3$, with absolute implied constant. This means furthermore that if $L > \alpha n^6$, for some absolute constant $\alpha > 0$, we have

$$H \gg \frac{1}{n} \frac{L}{\log L},$$

and that in consequence we then have the sieve estimate

$$\begin{aligned} & \mathbf{P}(\det(T - X_k) \text{ has small splitting field}) \\ & \leq (1 + 2nL^{n^2-n+1} \exp(-\eta_n k)) \frac{n \log L}{L}. \end{aligned}$$

We select

$$L = \left(\frac{1}{2n} \exp(k\eta_n) \right)^{1/n^2}.$$

If this quantity satisfies $L > \alpha n^6$, we can proceed to obtain the upper bound

$$\begin{aligned} & \mathbf{P}(\det(T - X_k) \text{ has small splitting field}) \\ & \ll n \exp\left(-\frac{k\eta_n}{n^2}\right) \frac{k\eta_n}{n^2} \ll k \exp\left(-\frac{k\eta_n}{n^2}\right) \end{aligned}$$

where the implied constant is absolute.

On the other hand, if $L \leq \alpha n^6$, such an estimate is trivial if the implied constant in question is sufficiently large, so that the Proposition is proved. \square

This statement suggests some fairly interesting questions. In general, there is a lot of interest in probability theory in phenomena exhibiting what is called ‘abrupt transition’, ‘phase transition’, ‘cut-off phenomenon,’ or ‘threshold phenomenon’ (see, e.g., [111, 3.3] for a discussion): in the context of random walks, it means intuitively that some event defined for a sequence $(X_{n,k})$ of walks on the groups G_n has the property that it happens with very small probability until some ‘time’ k_n , and happens with probability almost one a very short time after k_n . In the case of walks on finite groups, the ‘event’ is often simply the approximation to the uniform distribution; the most famous example is the analysis of card shuffling, seen as random walks on the permutation groups \mathfrak{S}_n : six ‘riffle shuffles’ do not mix a deck of cards well, but seven typically do (see [111, 3.2]).

Here we consider the sequence of walks $(X_{n,k})_{k \geq 0}$ on $SL(n, \mathbf{Z})$ with respect to the generators above, denoted S_n to emphasize the dependency on n ,¹⁵ and we look at the reducibility of $\det(T - X_{n,k})$, or the Galois group of its splitting field more generally. According to Proposition 7.23, taking into account the approximate value of η_n , this happens with exponentially small probability as soon as k is larger than, say, $Cn^5 \log n$, for some (absolute) constant $C > 0$. It seems therefore interesting to further investigate this transition; this may be phrased in various ways, involving the random variables

$$\begin{aligned}\tau_n &= \min\{k \geq 1 \mid \det(T - X_{n,k}) \text{ is irreducible}\}, \\ \tau_n^* &= \max\{k \geq 1 \mid \det(T - X_{n,k}) \text{ is reducible}\},\end{aligned}$$

which satisfy $\tau_n \leq \tau_n^* < +\infty$ (almost surely), and the variants with irreducibility replaced with having maximal Galois group. One may ask for information about the distribution of those random variables. From Proposition 7.23, since

$$\begin{aligned}\mathbf{P}(\tau_n = k + 1) &\leq \mathbf{P}(\tau_n^* = k) \leq \mathbf{P}(\det(T - X_{n,k}) \text{ is reducible}) \\ &\leq \max(1, Ck \exp(-k \frac{\eta_n}{n^2}))\end{aligned}$$

for some absolute constant $C \geq 0$, it follows easily that

$$\mathbf{E}(\tau_n) \ll n^5 \log n \quad (7.26)$$

for $n \geq 2$, where the implied constant is absolute.

For a first step, note that at the very least $\det(T - X_{n,k})$ is reducible for $k \leq t_n$ where t_n is the first (random) time when all basis vectors have been moved at least once. Since multiplying by $\xi_{n,k}$ involves moving one only of the n basis vectors, chosen uniformly, t_n is the stopping time for the ‘coupon collector problem’. Besides the obvious bound $t_n \geq n$, it is well known (see, e.g., [41, IX.3.d]) that

$$\mathbf{E}(t_n) = n(\log n + \gamma) + O(1), \text{ for } n \geq 1, \quad \mathbf{V}(t_n) \sim \zeta(2)n^2 \text{ as } n \rightarrow +\infty;$$

in fact, t_n is the sum of n independent random variables with geometric distribution with parameter $(n - k + 1)/n$, $1 \leq k \leq n$, each of which has expectation $n/(n - k + 1)$, hence the expectation of t_n is

$$n \sum_{1 \leq j \leq n} \frac{1}{j}.$$

¹⁵ Of course, this is where uniformity in n becomes a really interesting feature and well worth the effort. Also this discussion should be applicable to more general situations, with different generating sets, different groups, different ‘events’...

Table 7.1 *Transition times for random walks on $SL(n, \mathbf{Z})$*

n	Samples	Average of t_n	Average of τ_n	Ratio	Average of t_n/τ_n
10	100 000	29.258	38.452	1.314240	1.428859
15	100 000	49.824	62.785	1.260139	1.350101
20	100 000	71.916	88.371	1.228816	1.302885
25	70 000	95.112	115.366	1.212937	1.277969
30	70 000	119.900	143.387	1.195884	1.253686
40	70 000	171.154	201.536	1.177512	1.226495
50	35 000	225.101	263.028	1.168489	1.211726
75	30 000	367.688	422.558	1.149229	1.184520
100	30 000	519.610	590.741	1.136893	1.167134

Note: When $n = 100$ we stopped the walk after 1200 steps, and in 44 cases among the 30 000 samples, the time t_n had not yet been reached; hence the data is very slightly off.

The gap between the upper and lower bounds for the time to become irreducible with large probability is quite important, and it seems intuitively clear that the lower bound should be much closer to the truth. This is also suggested by numerical experiments, and in fact those suggest that the answer to the question in the following problem might well be ‘Yes’:

Problem 7.24 Does there exist a constant $c \geq 1$ such that

$$\mathbf{E}(\tau_n) \sim c\mathbf{E}(t_n) \sim cn \log n$$

for $n \rightarrow +\infty$? Is it in fact true with $c = 1$?

Numerically, we simulated the random walk (and the coupon collector problem involved in computing t_n) for $n = 10, 15, 20, 25, 30, 40, 50, 75, 100$, using PARI/GP and MAGMA. The data we obtained are in Table 7.1, where the first column is the number of samples used (e.g., 100 000 random walks with $n = 10$ were performed, and the empirical average of t_n was 29.26549, the empirical average of τ_n was 38.53915).

It seems very difficult to improve either the upper bound (7.26) for $\mathbf{E}(\tau_n)$, or the upper bound in Proposition 7.23. In particular, it is known that the order of magnitude of Kassabov’s estimate of the Kazhdan constant ε_n for the generating sets S_n is optimal (Zuk has pointed out that this constant must be $\geq \sqrt{2/n}$, see [118, p. 149]). So this ‘arithmetic transition’ problem for random walks looks very challenging.

On the other hand, it is possible to change the generating set and improve the Kazhdan constant. Indeed, Hadad has proved [52], developing further the methods of Shalom and Kassabov, that there exist a constant k and generating sets S_n for $SL(n, \mathbf{Z})$, $n \geq 3$, such that $|S_n| \leq k$ for all n , and the associated

Kazhdan constant κ_n is uniformly bounded away from zero for $n \geq 3$: $\kappa_n \gg 1$ for $n \geq 3$. (A result of this type should be compared with the effect of the Riemann Hypothesis over finite fields, see Corollary 8.10 in the next chapter). This means, using the same argument as in the proof of Proposition 7.23, that the corresponding random walks $(X_{n,k})$ satisfy

$$\mathbf{P}(\det(T - X_{n,k}) \text{ is reducible}) \ll k \exp(-\delta k/n^2)$$

for all $n \geq 3$ and $k \geq 1$, where $\delta > 0$ is some absolute constant. In particular, the expectation of the analogue of the transition time τ_n is now bounded by (a constant times) $n^2 \log n$.

7.8 Other groups

We have concentrated our attention on applications of the sieve to some specific arithmetic groups. This is partly because of the convenience and concreteness arising from such a choice, partly because those groups suggested natural problems that were both appealing and accessible to the sieve technique. This is not to suggest that Theorem 7.4 is the only application of Proposition 7.2 to obtain sieves for random walks on finitely generated groups.

Indeed, if we simply assume that the generating set contains 1, to avoid any issue with periodicity, then two conditions are essential for a group G to be susceptible to sieve applications as described in this chapter:

- It should have many finite quotients.
- It should satisfy Property (τ) with respect to some family of interesting quotients, or even Property (T) .

The first condition means that, essentially, we can or should assume that G is *residually finite*, i.e., for any $g \in G$ not trivial, there exists a finite group H and a homomorphism $f : G \rightarrow H$ such that $f(g) \neq 1$. There is an abundance of such groups – for instance, any finitely generated subgroup of a linear group $GL(n, \mathbf{C})$ is residually finite (a theorem of Mal'cev, see [57, III.18, III.20] for references).

The second condition seems more restrictive. However, there are now many examples of groups which are known to have Property (T) ; for instance, Shalom has recently shown that $SL(n, \mathbf{Z}[x_1, \dots, x_m])$ has this property for any $n \geq 1$ if $m \geq n + 3$; see his ICM paper [119] for other results and references. Moreover, Property (T) is, in a sense, a ‘generic’ property: Silberman [121] has in fact proved that certain types of finitely generated groups defined by random presentations have Property (T) with very high probability. In addition, there is much ongoing interest (and success) in finding examples of groups with Property (τ) .

For instance, following a breakthrough of Helfgott [59], work of Bourgain and Gamburd and Bourgain, Gamburd and Sarnak [13, 15] has shown that any discrete subgroup $\Gamma \subset SL(2, \mathbf{Z})$ which is Zariski-dense in $SL(2, \mathbf{Z})$ has Property (τ) with respect to the ‘congruence’ subgroups $\text{Ker}(\Gamma \rightarrow SL(2, \mathbf{Z}/q\mathbf{Z}))$, where q is any squarefree number. Moreover, it may be hoped that this will be generalized to, e.g., $\Gamma \subset SL(n, \mathbf{Z})$ which would lead to very general sieve settings.

Also, even if the group G of interest does not (or is not known to have) Property (T/τ) , there might still be applications of sieve. We saw this in Chapter 6, with an abelian group $G = \mathbf{Z}$, and in this chapter with the geometric applications involving the mapping class groups. Whether the mapping class group of a closed surface Σ_g of genus $g \geq 2$ has Property (T) was an outstanding question; for $g = 2$, Taherkhani [128] proved that the mapping class group does not have Property (T) , by a direct computation of a finite index subgroup with infinite abelianization; and a recent paper of Andersen [5] announces that this is the case for all $g \geq 2$, using very different ideas. It is not yet known whether Property (τ) holds, e.g., for finite index subgroups (the representations that Andersen uses are faithful, hence do not factor through such a subgroup).

If we were to consider such an abstract residually finite finitely generated group Γ with Property (τ) for finite index subgroups, one may object that in general there is no ‘natural’ family of maps $(G \rightarrow G_\ell)$ which is a good candidate to complete the sieve setting; the family of reduction maps modulo primes which was used previously makes no sense in general, and taking an overly large family (e.g., all surjective homomorphisms to finite groups) is unlikely to be of use because the linear disjointness property (Definition 2.16) that encapsulates the desired independence of the various maps will not hold. So we want to point out a related family (ρ_ℓ) that *does* satisfy this condition.

Let $\tilde{\Lambda}$ be the set of surjective homomorphisms

$$\rho : G \rightarrow \rho(G) = H$$

where H is a non-abelian finite simple group, and let $\Lambda \subset \tilde{\Lambda}$ be a set of representatives for the equivalence relation defined by $\rho_1 \sim \rho_2$ if and only if there exists an isomorphism $\rho_1(G) \rightarrow \rho_2(G)$ such that the triangle

$$\begin{array}{ccc} & G & \\ \rho_1 \swarrow & & \searrow \rho_2 \\ \rho_1(G) & \xrightarrow{\quad} & \rho_2(G) \end{array} \quad (7.27)$$

commutes.

Lemma 7.25 *The system $(\rho)_{\rho \in \Lambda}$ constructed in this manner is linearly disjoint.*

This is an easy adaptation of classical variants of the Goursat–Ribet lemmas, and is left as an exercise (see, e.g. [107, Lemma 3.3] and Lemma 3.7 in [34], where a result of this type is attributed to P. Hall).

The next necessary step in an application of sieve, in practice, would be to gain some knowledge of Λ , and in particular one would probably need to know something of the distribution of the orders of the finite simple quotient groups of G which occur as targets (with the goal, maybe, of using the sieve support $\{\rho \in \Lambda \mid |\rho(G)| \leq L\}$ for some L). This type of question is of course in itself an interesting problem, and in fact deserving of book-length treatment (see, e.g., [92]).

8

Sieving for Frobenius over finite fields

In this final chapter, we will describe the use of the large sieve to study the average distribution of (geometric) Frobenius conjugacy classes in Galois groups of coverings of algebraic varieties over finite fields, or equivalently in a more geometric language that we will use instead, in finite monodromy groups of sheaves obtained by reduction of integral ℓ -adic sheaves. This sieve is a good example (in fact, the most interesting at the moment) of a coset (conjugacy) sieve, as defined in Section 3.3.

This type of sieve was introduced in [80], and its strengthening was the motivation for the paper from which this book evolved. We will recall enough of the previous work to make the argument independent of results in [80].

As explained in Example 4.10, there is nothing to prevent adapting the ideas to sieve for Frobenius conjugacy classes over number fields, except that really good results depend at present on assuming some form of the Generalized Riemann Hypothesis (though weaker unconditional bounds are possible, see D. Zywinia's preprint 'The large sieve and Galois representations', 2007).

Contrary to what we have done in all previous applications of the sieve, we have not attempted to give entirely self-contained *definitions*; here, we need to introduce some 'black boxes'. Hopefully, the examples of applications (which we can, and do, describe from scratch) will be sufficiently interesting to encourage interested readers to get better acquainted with the foundations and in particular with Deligne's work on the Riemann Hypothesis over finite fields. In fact, we start by a section explaining the problem of Katz that was solved qualitatively by Chavdarov and which is the motivating problem for [80]. After this, we will describe the general, more abstract, setting of the sieve for Frobenius, before going on to prove the basic sieve statements (refining somewhat what was done in [80]), and then considering applications, old and new.

8.1 A problem about zeta functions of curves over finite fields

The motivating problem is a question of Katz which was first considered by Chavdarov in [22]. This story starts with a very concrete and classical diophantine question: what is the number of solutions of a system of polynomial equations over a finite field? More precisely, we concentrate on curves over finite fields, and let C/\mathbf{F}_q denote such a curve, which we assume to be a smooth, projective, geometrically irreducible curve. If the reader is unsure of the precise meaning of this ('curve' should have an intuitive meaning . . .), it is possible to restrict attention in this section to *plane* curves of this type. Then giving C amounts to giving a homogeneous polynomial $f \in \mathbf{F}_q[X, Y, Z]$, which is irreducible as an element of $\bar{\mathbf{F}}_q[X, Y, Z]$ (this is what 'geometrically irreducible' means; in general, the adjective 'geometric' relates to notions defined over an algebraically closed field containing the base field), and which is non-singular in the sense that $(0, 0, 0)$ is the only solution in $\bar{\mathbf{F}}_q^3$ of the set of equations

$$\begin{cases} f(x, y, z) = 0, \\ \partial_x f(x, y, z) = 0, \\ \partial_y f(x, y, z) = 0, \\ \partial_z f(x, y, z) = 0 \end{cases}$$

(in most cases, a 'randomly chosen' polynomial will work). If we are given such a polynomial f , the associated curve is the set of non-zero solutions $(x, y, z) \in \bar{\mathbf{F}}_q^3$ to $f(x, y, z) = 0$, except that because of the homogeneity of the equations, we take equivalence classes for the relation

$$(x, y, z) \sim (\lambda x, \lambda y, \lambda z)$$

for any $\lambda \in \bar{\mathbf{F}}_q^\times$ (this homogeneity is the reason we interpret the set of solutions as one-dimensional: there are three variables, one equation $f = 0$, so we would expect solutions depending 'on two parameters', and once the 'obvious' one, namely λ , is removed, there is – or should be – a single parameter left, which is the intuitive description of a curve). The curve is, informally, 'given by the equation $f = 0$ in the projective plane'.

Given the curve C (or the polynomial f), we are interested in the set of rational points $C(\mathbf{F}_q)$, which for a plane curve means a restriction to those $(x, y, z) \in \mathbf{F}_q^3$ which satisfy $f(x, y, z) = 0$, up to the above equivalence with $\lambda \in \mathbf{F}_q^\times$. More generally, since for any n there is an extension field \mathbf{F}_{q^n} , we can look at $C(\mathbf{F}_{q^n})$, the set of solutions with coordinates in this field. The first important fact is that all these sets are finite, since $\mathbf{F}_{q^n}^3$ itself is finite.

The basic diophantine issue is then to understand the sets $C(\mathbf{F}_{q^n})$ for fixed C and $n \geq 1$ arbitrary. The first question is to understand the number of solutions (further questions may of course be raised, but this seems the most basic), and for this purpose, it is natural to look at the *generating function* associated to the sequence $|C(\mathbf{F}_{q^n})|$, namely the formal power series

$$\sum_{n \geq 1} |C(\mathbf{F}_{q^n})| T^n \in \mathbf{Z}[[T]] \quad (8.1)$$

or rather, since it turns out to be better behaved (and closely analogous to the classical Riemann zeta function, although this is not clear from this formal definition . . .), the so-called *zeta function of C* , defined as the equally formal power series

$$Z(C, T) = \exp \left(\sum_{n \geq 1} |C(\mathbf{F}_{q^n})| \frac{T^n}{n} \right). \quad (8.2)$$

Notice that (8.1) is the logarithmic derivative of this power series; since $Z(C, 0) = 1$, there is the same amount of information in $Z(C, T)$ or in its logarithmic derivative.

A remarkable result states that this formal power series is quite special.

Theorem 8.1

(1) *The formal power series $Z(C, T)$ represents a rational function of the type*

$$Z(C, T) = \frac{P(T)}{(1 - T)(1 - qT)}$$

where $P \in \mathbf{Z}[T]$ is a polynomial with integer coefficients of some degree $2g$, $g \geq 0$, such that $P(0) = 1$ and

$$q^g T^{2g} P\left(\frac{1}{qT}\right) = P(T). \quad (8.3)$$

(2) *All complex zeros α of P satisfy $|\alpha| = 1/\sqrt{q}$.*

The first part is due to F. K. Schmidt in this generality, and the second, which is the analogue of the Riemann Hypothesis in this case, to A. Weil; in both cases, earlier results were known, due in particular to E. Artin, and H. Hasse. The integer g in the theorem, which only depends on the (geometric) curve, is the *genus* of the curve.

Example 8.2 The first part of this theorem may be phrased equivalently as follows: factor

$$P(T) = \prod_{1 \leq i \leq 2g} (1 - \alpha_i T),$$

then comparing the power series expansion of the logarithmic derivative of $Z(C, T)$ (namely, (8.1)) and of its representation as

$$\frac{(1 - \alpha_1 T) \cdots (1 - \alpha_{2g} T)}{(1 - T)(1 - qT)},$$

we obtain the formula

$$|C(\mathbf{F}_{q^n})| = q^n + 1 - \sum_{1 \leq i \leq 2g} \alpha_i^n$$

for $n \geq 1$. Or, more concretely yet, the sequence $u_n = |C(\mathbf{F}_{q^n})|$ for $n \geq 1$ satisfies a linear recurrence of order $2g + 2$, since its generating series is itself a rational function:

$$\sum_{n \geq 1} u_n T^n = \frac{T}{1 - T} + \frac{qT}{1 - qT} - \sum_{1 \leq i \leq 2g} \frac{\alpha_i T}{1 - \alpha_i T},$$

so that multiplying through by a common denominator $P(T)(1 - T)(1 - qT)$ and stating that the coefficients of degree $> 2g + 2$ vanish, we obtain such a linear recurrence. One tends to see this as standard nowadays, but it is quite amazing: it means that knowing the number of points of C in the first few small fields \mathbf{F}_{q^n} is sufficient to know all the others; in particular, this gives the number of points in fields which intersect those small fields only in \mathbf{F}_q itself (e.g., those with large prime degree); further, note that C may well have no point at all with coefficients in $\mathbf{F}_q \dots$

The polynomials P which occur as the numerator of the zeta function of algebraic curves C/\mathbf{F}_q are the concrete subject of the Katz–Chavdarov problem, which informally may be stated as follows:

Problem 8.3 In what ways is P similar to a ‘random’ polynomial? In particular, is its factorization in $\mathbf{Q}[X]$, or its splitting field, typically the same as that of a ‘random’ polynomial?

This is a natural question. To make it more precise, we can see from Theorem 8.1, part (1), that we can not expect P to be really ‘generic’ (in the sense of Gallagher’s Theorem) because the equation (8.3) forces some relations among the roots of P , hence imposes a restriction on the Galois group of the splitting field of P . This is similar to the relation (7.17) for characteristic polynomials of symplectic matrices, and indeed, as we will recall, this is no coincidence.

Precisely, if γ is a root of P , then $(q\gamma)^{-1}$ is also a root. In terms of the parameters α_i introduced in Example 8.2, which are the inverses of the roots of

P , this means that for any i , $1 \leq i \leq 2g$, q/α_i is also among the α_i (it could be equal to α_i itself, if $\alpha_i = \pm\sqrt{q}$, but this is a rare occurrence).

In terms of Galois groups, this implies that it is possible to number the roots of P in pairs $(\alpha_{2i-1}, \alpha_{2i})_{1 \leq i \leq g}$, so that as a permutation group acting on $\{1, \dots, 2g\}$ used as labels for the roots, the Galois group of P is contained inside the subgroup W_{2g} of signed permutations; recall these are the permutations σ such that the g pairs $(2i-1, 2i)$ are themselves permuted.¹ We have $|W_{2g}| = 2^g g!$, and there is an exact sequence

$$1 \rightarrow \{\pm 1\}^g \rightarrow W_{2g} \xrightarrow{P} \mathfrak{S}_g \rightarrow 1. \quad (8.4)$$

So a more precise form, still informal, of the question of Katz is:

Problem 8.4 How does the splitting field of P compare with the splitting field of a typical polynomial for which (8.3) holds? Is the Galois group typically isomorphic to W_{2g} ?

We will say that a polynomial $P \in \mathbf{Z}[X]$ which is of degree $2g$ and satisfies $P(0) = 1$ and (8.3) is *q-symplectic* (for reasons which will become clear soon). Now it is a fact that most *q*-symplectic polynomials have maximal Galois group, namely W_{2g} , in the same sense that most monic polynomials of degree r have Galois group the full symmetric group (see [27] for $q = 1$, or [80, Remark 7.4]). This is in any case the natural guess.

To finally make precise the question that will be solved in the next sections, a last decision must be taken: which sets of curves are we going to select in which to look for polynomials P with maximal splitting field? (Of course, for a given C , the answer may very well turn out to be something different, just as there are polynomials with Galois group different from the symmetric group.) The choice of Katz and Chavdarov is to look at an *algebraic family* of curves. Indeed, this is how the powerful methods developed by Grothendieck, Deligne, and others, will be most effective.

Formally, such a family of curves is the data of a (surjective) morphism $\mathcal{C} \xrightarrow{\pi} U$ of algebraic varieties over \mathbf{F}_q such that U (the ‘parameter variety’) is typically an open set in some affine space, and all the fibers of π are themselves smooth, geometrically irreducible, projective curves, with fixed genus g . In keeping with the elementary viewpoint of this section, the reader may see this

¹ It is not necessarily the case that $\alpha_{2i-1}\alpha_{2i} = q$, because of the exceptional case of $\pm\sqrt{q}$, but the number of such exceptions is even (since the degree is $2g$), and they occur either as integers of the form $\pm\sqrt{q}$, if q is a square, in which case they can be paired arbitrarily; or as pairs $(\sqrt{q}, -\sqrt{q})$ if q is not a square. In both cases, the property described still holds. Moreover, these exceptions can not occur if $g \geq 2$ and P is irreducible.

concretely (in a special case) as the data of a polynomial $f \in \mathbf{F}_q[X, Y, Z, T]$, such that for all $n \geq 1$ and all $t \in \mathbf{F}_{q^n}$ (with finitely many exceptions maybe), the specialized polynomial $f_t = f(X, Y, Z, t) \in \mathbf{F}_{q^n}[X, Y, Z]$ defines a smooth, geometrically irreducible, projective curve of genus g over \mathbf{F}_{q^n} , with equation

$$C_t : f(x, y, z, t) = 0$$

(where t is fixed). Fixing $n = 1$, this gives a set of curves over \mathbf{F}_q , which usually contains roughly q elements, and a corresponding set of polynomials P_t from the numerators of the zeta functions of C_t . Here is then the precise question:

Problem 8.5 Let $\mathcal{C} \rightarrow U$ be an algebraic family of smooth, geometrically irreducible, projective curves of genus g over \mathbf{F}_q . For $t \in U(\mathbf{F}_q)$, how often is the Galois group of the splitting field of P_t as large as possible, i.e., isomorphic to W_{2g} ?

It turns out that the answer is ‘most of the time’, in a quantitative way, but *only* under a further assumption on the family of curves. This condition is something which does not arise in the case of ‘all’ polynomials of bounded height, and although it is typically expected to hold, it is not easy to check. We now give the most concrete example where it is known, and where our results will apply.

Example 8.6 Let $g \geq 1$ be an integer, and let $f \in \mathbf{F}_q[X]$ be a monic polynomial of degree $2g$ which is squarefree (i.e., does not have multiple roots in $\bar{\mathbf{F}}_q$).² Consider the polynomial

$$h = Y^2 - f(X)(X - T) \in \mathbf{F}_q[X, Y, T]$$

and its homogenized version

$$\tilde{h} = Z^{2g+1} h\left(\frac{X}{Z}, \frac{Y}{Z}, \frac{T}{Z}\right) \in \mathbf{F}_q[X, Y, Z, T].$$

For any $t \in \bar{\mathbf{F}}_q$ which is not a root of f , it is a standard fact of algebraic geometry that the curve with affine equation

$$h(x, y, t) = 0$$

is a smooth geometrically irreducible *affine* curve; however, for $g \geq 2$, the projective curve $\tilde{h}(x, y, z, t) = 0$ is not smooth (there is a problem ‘at infinity’, i.e., when $z = 0$, where the only solutions are $(0, y, 0) \sim (0, 1, 0)$, and the partial derivatives also vanish at this point if $g \geq 2$). There is another standard

² This polynomial need not satisfy (8.3), it plays a different role . . .

technique of algebraic geometry to remove the singularity while keeping the affine part of the curve unchanged. As is customary, we will simply call this ‘the smooth projective model of the affine curve $y^2 = f(x)(x - t)$ ’. Then, for a fixed polynomial f , this will provide, by varying t among elements where $f(t) \neq 0$, our standard family of numerators of zeta functions, which are q -symplectic polynomials of degree $2g$. We will show, following and slightly improving [80], that for most values of t , the splitting field of this polynomial is as large as possible; see Theorem 8.15.

Note that curves given by equations of the type above are examples of what are called *hyperelliptic curves*. When $g = 1$, they are elliptic curves. For basic information on such curves, including a more intrinsic definition and its link with the equations above, see, e.g. [90, 7.4.3].

To detect W_{2g} as a subgroup of \mathfrak{S}_{2g} , we will use the result in the following exercise.

Exercise 8.1 Let $g \geq 1$ be an integer; consider W_{2g} as a subgroup of \mathfrak{S}_{2g} , and let $G \subset W_{2g}$ be a subgroup such that: (i) G contains a transposition and acts transitively on $\{1, \dots, 2g\}$; (ii) if $g \geq 2$, the projection $p(G) \subset \mathfrak{S}_g$ (see the exact sequence (8.4)) contains a transposition and an m -cycle for some prime $m > g/2$. Show that $G = W_{2g}$. [Hint: Use the lemma of Gallagher quoted in the proof of Theorem 4.2, or see [80, Lemma 7.1].]

8.2 The formal setting of the sieve for Frobenius

We now describe the precise setting of the sieve that is suitable for solving Problem 8.5. Here we will use the language of algebraic geometry and ℓ -adic sheaves without hesitation. The link with the problem itself will be explained in Section 8.6. For basic references, we refer to [97], [72], [77, Chapters 9, 10].

Let q be a power of a prime p , and let U/\mathbf{F}_q be a smooth affine geometrically connected algebraic variety of dimension $d \geq 1$ over \mathbf{F}_q (which, as usual, is of characteristic p). Put $\bar{U} = U \times \bar{\mathbf{F}}_q$, the extension of scalars to an algebraic closure of \mathbf{F}_q .

Let $\bar{\eta}$ denote a geometric generic point of U . We then have at our disposal two profinite groups: the arithmetic fundamental group $\pi_1(U, \bar{\eta})$, and the geometric fundamental group $\pi_1(\bar{U}, \bar{\eta})$. They sit in an exact sequence

$$1 \rightarrow \pi_1(\bar{U}, \bar{\eta}) \rightarrow \pi_1(U, \bar{\eta}) \xrightarrow{d} \hat{\mathbf{Z}} \rightarrow 1,$$

where the last map is called the *degree*. These groups are analogues both of the (absolute) Galois group of a field – and they may indeed be interpreted as such in many cases – and of the topological fundamental group of a topological space (as described in Appendix H, for instance). In particular, to give a surjective homomorphism $\pi_1(U, \bar{\eta}) \rightarrow H$, for H a finite group, is equivalent with giving an étale Galois covering $U_H \rightarrow U$ with Galois group H , and similarly with $\pi_1(\bar{U}, \bar{\eta})$. The distinction between the arithmetic and geometric fundamental group arises from the fact that there are étale coverings which are ‘geometrically’ trivial, namely those given by extensions of scalars

$$U \times_{\mathbf{F}_q} \mathbf{F}_{q^n} \rightarrow U$$

which become trivial when extended to $\bar{\mathbf{F}}_q$; not coincidentally, these coverings have cyclic Galois groups, hence the occurrence of $\hat{\mathbf{Z}}$ as quotient of the arithmetic fundamental group modulo the geometric one (concretely, $\hat{\mathbf{Z}}$, the profinite completion of \mathbf{Z} , can be seen as the profinite group which admits for every $d \geq 1$ a unique finite index subgroup with quotient $\mathbf{Z}/d\mathbf{Z}$).

For any $n \geq 1$ and any rational point $x \in U(\mathbf{F}_{q^n})$, there is (by definition of rational points!) a morphism $\text{Spec } \mathbf{F}_{q^n} \rightarrow U$ which ‘is’ x , and hence by functoriality, an induced map from the fundamental group of $\text{Spec } \mathbf{F}_{q^n}$ (which is isomorphic to the Galois group of \mathbf{F}_{q^n} , and hence topologically generated by the n -th power of the arithmetic Frobenius $x \mapsto x^q$, or of its inverse the geometric Frobenius automorphism) to the fundamental group of U . The conjugacy class of the image of the geometric Frobenius automorphism is well-defined, and is called the geometric Frobenius at x . We denote it by Fr_{x,q^n} , or simply Fr_x when the field of definition of x is clearly fixed (note that seeing x as defined over a larger field changes the Frobenius automorphism). In the exact sequence above, we have $d(\text{Fr}_{x,q^n}) = -n$ for all n (the minus sign comes from taking the geometric Frobenius; its inverse, the arithmetic Frobenius, has degree n).

Since we will be interested in the behaviour of Fr_{x,q^n} for fixed n , as reflecting interesting arithmetic properties of the rational point x , we see that we are exactly in the situation of the (conjugacy) coset sieve of Section 3.3 with

$$G = \pi_1(U, \bar{\eta}), \quad G^g = \pi_1(\bar{U}, \bar{\eta}), \quad G/G^g \simeq \text{Gal}(\bar{\mathbf{F}}_q/\mathbf{F}_q) \simeq \hat{\mathbf{Z}}. \quad (8.5)$$

We then naturally will take the siftable set to be $X = U(\mathbf{F}_{q^n})$ with the map $x \mapsto F_x = \text{Fr}_{x,q^n}$, which is a conjugacy class such that $d(F_x) = -n$ is fixed for all $x \in X$. In fact, we will simply assume $n = 1$ for simplicity; any result we obtain, as long as the dependency on the base field is explicit, can then be applied to all extensions \mathbf{F}_{q^n} by replacing U by its extension of scalars to \mathbf{F}_{q^n} .

To wrap up the sieve setting, we need surjective maps from G to finite groups. There is an abundance of those, since G is profinite, and they correspond to

Galois étale coverings of U , as we already mentioned. Partly for reasons of convenience, and partly because the natural examples are of this type, we assume given a family of homomorphisms

$$\rho_\ell : \pi_1(U, \bar{\eta}) \rightarrow GL(r, k_\ell)$$

which are continuous in the sense that $\text{Ker } \rho_\ell$ is closed in $\pi_1(U, \bar{\eta})$, for ℓ in a subset Λ of the set of prime numbers (different from p), and where k_ℓ is a finite field of characteristic ℓ while r is independent of ℓ .

We do not assume that ρ_ℓ is onto (this will rarely be the case), but instead we define the *arithmetic monodromy group* of ρ_ℓ by

$$G_\ell = \text{Im}(\rho_\ell) = \rho_\ell(G),$$

and the *geometric monodromy group* by

$$G_\ell^g = \rho_\ell(G^g),$$

so that we have our formal coset sieve setting

$$\Psi = (Y, \Lambda, (\rho_\ell : Y \rightarrow Y_\ell)),$$

where $Y \subset G^\sharp$ (respectively $Y_\ell \subset G_\ell^\sharp$) is the set of those conjugacy classes g^\sharp such that $d(g^\sharp) = -1$.

Note that, in geometric terms, this family of homomorphisms corresponds to a family of étale Galois coverings $U_\ell \rightarrow U$ with Galois group G_ℓ , a subgroup of $GL(r, k_\ell)$. A standard terminology, arising from another equivalent interpretation of those coverings, is that ρ_ℓ is a *lisse sheaf of k_ℓ -vector spaces*. When using the sheaf-theoretic language,³ it is customary to use curly letters \mathcal{F}_ℓ instead of ρ_ℓ to denote a sheaf. The k_ℓ -vector space on which ρ_ℓ acts can be naturally identified, in sheaf terms, as the fiber $(\mathcal{F}_\ell)_{\bar{\eta}}$ of the sheaf over the generic point.

In the next section, we require in some cases that the system of k_ℓ -adic sheaves is obtained by reduction from a compatible system of integral ℓ -adic sheaves. We review the definition.

Definition 8.7 *Let U/\mathbb{F}_q be as above. A system (ρ_ℓ) of continuous homomorphisms $\pi_1(U, \bar{\eta}) \rightarrow GL(r, k_\ell)$ for ℓ in a subset Λ of primes distinct from p , is a compatible system if there exists a number field K/\mathbb{Q} with ring of integers \mathbf{Z}_K , and for all $\ell \in \Lambda$ a prime ideal $\lambda \subset \mathbf{Z}_K$ above ℓ with residue field $\mathbf{Z}_K/\lambda \simeq k_\ell$, and a continuous homomorphism*

$$\tilde{\rho}_\ell : \pi_1(U, \bar{\eta}) \rightarrow GL(r, \mathbf{Z}_\lambda),$$

³ Which it is not really necessary to know precisely here.

where \mathbf{Z}_λ is the ring of integers in the completion K_λ of K at λ , such that:

- for every ℓ , the reduction of $\tilde{\rho}$ modulo $\lambda\mathbf{Z}_\lambda$ is isomorphic to ρ_ℓ ;
- for every ℓ , every extension field \mathbf{F}_{q^n} of \mathbf{F}_q , every $u \in U(\mathbf{F}_{q^n})$, the reversed characteristic polynomial

$$\det(1 - T \tilde{\rho}_\ell(\text{Fr}_{u, q^n})) \in \mathbf{Z}_\lambda[T]$$

has coefficients in the ring of integers \mathbf{Z}_K of K , and is independent of ℓ .

In sheaf-theoretic language, where ρ_ℓ corresponds to a lisse sheaf \mathcal{F} of k_ℓ -vector spaces, we say instead that there are étale sheaves of free \mathbf{Z}_λ -modules $\tilde{\mathcal{F}}_\ell$ such that $\mathcal{F}_\ell = \tilde{\mathcal{F}}_\ell / \lambda \tilde{\mathcal{F}}_\ell$ for all ℓ .

A consequence of having a compatible system that we will use is the following. Denote by $|U|$ the set of *closed points* of U , which can be identified with the set of orbits under the action of $\text{Gal}(\bar{\mathbf{F}}_q/\mathbf{F}_q)$ of the points $x \in U(\bar{\mathbf{F}}_q)$. For $x \in |U|$, define the (geometric) Frobenius conjugacy class Fr_x to be $\text{Fr}_{y, \deg(x)}$, for any rational point y in the orbit which is x (the resulting conjugacy class is well-defined in $\pi_1(U, \bar{\eta})$), the degree $\deg(x)$ of the closed point being the cardinality of the orbit, or in other words the degree of the residue field at x , so that $y \in U(\mathbf{F}_{q^{\deg(x)}})$ for all y in the orbit. Then define the L -function of the system to be the formal power series

$$L(T) = \prod_{x \in |U|} \det(1 - T \tilde{\rho}_\ell(\text{Fr}_x))^{-1},$$

which, according to Definition 8.7, is a power series with coefficients in \mathbf{Z}_K , and is independent of the choice of ℓ . As proved by Grothendieck, $L(T)$ is a rational function in $K(T)$, and hence the *opposite of its degree*, defined as the difference of the degree of the denominator and that of the numerator, is an integer independent of ℓ , which is the *Euler–Poincaré characteristic* of the compatible system (with compact support). It does not depend on the base field \mathbf{F}_q either, but only on the extension of scalars \bar{U} , and the system of representations, or sheaves, on \bar{U} , i.e., on the maps

$$\pi_1(\bar{U}, \bar{\eta}) \rightarrow GL(r, \mathbf{Z}_\lambda).$$

If we take the trivial compatible system ($r = 1$ and $\rho_\ell = 1$ for all ℓ), we obtain the zeta function of U as L -function, and the Euler–Poincaré characteristic is called simply the Euler–Poincaré characteristic of \bar{U} . (From Theorem 8.1, we can see that the Euler–Poincaré characteristic for a smooth projective geometrically irreducible curve of genus g is $2 - 2g$). In Section 8.4, we will recall the original definition by means of alternating sums of Betti numbers.

8.3 Bounds for sieve exponential sums

Having set up precisely the sieve for Frobenius in the previous section, we now turn to the problem of estimating the large sieve constant, which we approach by means of exponential sums, as in Proposition 3.7. By definition (see (3.14)), those sums are of the following type:

$$W(\pi, \tau) = \frac{1}{\sqrt{|\hat{\Gamma}_m^\pi| |\hat{\Gamma}_n^\tau|}} \sum_{u \in U(\mathbb{F}_q)} \text{Tr}(\pi(\rho_m(\text{Fr}_{u,q}))) \overline{\text{Tr}(\tau(\rho_n(\text{Fr}_{u,q})))},$$

where $m, n \in S(\Lambda)$ are squarefree integers divisible only by primes $\ell \in \Lambda$, and π (respectively τ) is an irreducible representation of G_m (respectively G_n).

As in [80], we will prove two bounds for these sums, one of which requires some assumptions on the ramification of the maps ρ_ℓ , while the other asks that the sheaves (ρ_ℓ) form a *compatible system*, and is restricted to one-parameter families.⁴ Moreover, it is necessary in both cases to assume that the system is *linearly disjoint* (this will be quite transparent in the proof).

We recall that for $m \in S(\Lambda)$, Π_m is a set of representatives of irreducible representations of G_m for the equivalence relation of isomorphism restricted to G_g^g , with $1 \in \Pi_m$, and that Π_m^* is the subset of Π_m determined by the condition that when writing $\pi = \boxtimes \pi_\ell$ with π_ℓ a representation of G_ℓ , we have $\pi_\ell \neq 1$ for all $\ell | m$, and in addition the function $\text{Tr } \pi(x)$ is not identically zero for $x \in Y_\ell$.

Proposition 8.8 *Assume that the representations (ρ_m) for $m \in S(\Lambda)$ are such that, for all squarefree numbers m divisible only by primes in Λ , the map*

$$\pi_1(\overline{U}, \bar{\eta}) \rightarrow G_m^g = \prod_{\ell|m} G_\ell^g \quad (8.6)$$

*is onto.*⁵ *With notation as before and as in Proposition 3.7, we have:*

(1) *If G_ℓ is a group of order prime to p for all $\ell \in \Lambda$, then*

$$W(\pi, \tau) = \delta((m, \pi), (n, \tau)) q^d + O\left(q^{d-1/2} |G_{[m,n]}| (\dim \pi)(\dim \tau)\right)$$

for $m, n \in S(\Lambda)$, $\pi \in \Pi_m^$, $\tau \in \Pi_n^*$, where the implied constant depends only on \overline{U} .*

⁴ Essentially because the analysis of wild ramification for ℓ -adic sheaves in higher dimensions seems not yet sufficiently developed to argue in general as we can do for curves.

⁵ It is easy to check that this condition is equivalent with the linear disjointness of Definition 2.16, namely that $Y \rightarrow Y_m$ is onto for any $m \in S(\Lambda)$, but it is more natural to state it in this manner, as it is the way it occurs in the proofs.

(2) If $d = 1$ (U is a curve) and if the sheaves \mathcal{F}_ℓ are of the form $\mathcal{F}_\ell = \tilde{\mathcal{F}}_\ell / \ell \tilde{\mathcal{F}}_\ell$ for some compatible system of torsion-free \mathbf{Z}_ℓ -adic sheaves $\tilde{\mathcal{F}}_\ell$, then

$$W(\pi, \tau) = \delta((m, \pi), (n, \tau))q + O\left(q^{1/2}(\dim \pi)(\dim \tau)\right)$$

where the implied constant depends only on the compactly-supported Euler–Poincaré characteristics of \overline{U} and of the compatible system $(\tilde{\mathcal{F}}_\ell)$ on \overline{U} .

Proof This is a generalization of Proposition 5.1 of [80], which corresponds to the case $m, n \in \Lambda$.

By (3.9) we can write

$$W(\pi, \tau) = \frac{1}{\sqrt{|\hat{\Gamma}_m^\pi| |\hat{\Gamma}_n^\tau|}} \sum_{u \in U(\mathbf{F}_q)} \text{Tr}([\pi, \bar{\tau}] \rho_{[m, n]}(\text{Fr}_{u, q}))$$

where the sum is the sum of local traces of Frobenius for a continuous representation

$$\pi_1(U, \bar{\eta}) \rightarrow GL((\dim \pi)(\dim \tau), \mathbf{C}).$$

The first step is to observe that, because this representation factors through the finite group $G_{[m, n]}$, it amounts to a finite-dimensional representation of a finite group in a field of characteristic zero, and as such it can be realized over a number field, and a fortiori over any other algebraically closed field of characteristic zero (see, e.g., [115, Section 12.3]). This means we can view $[\pi, \bar{\tau}]$ as a representation

$$\pi_1(U, \bar{\eta}) \rightarrow GL((\dim \pi)(\dim \tau), \overline{\mathbf{Q}}_\ell),$$

where ℓ is any prime number different from p . (But ℓ need not have anything to do with the primes dividing $[m, n]$.) Then we see that $W(\pi, \tau)$ is – up to a factor – the sum of local traces of Frobenius at points in $U(\mathbf{F}_q)$, acting on a finite-dimensional $\overline{\mathbf{Q}}_\ell$ -vector space, or to use sheaf language, acting on some lisse $\overline{\mathbf{Q}}_\ell$ -adic sheaf $\mathcal{W}(\pi, \tau)$ on U .

By the Grothendieck–Lefschetz Trace Formula, for *any* lisse sheaf \mathcal{F} of $\overline{\mathbf{Q}}_\ell$ vector spaces with $\ell \neq p$,⁶ the sum

$$\sum_{u \in U(\mathbf{F}_q)} \text{Tr}(\text{Fr}_{u, q} | \mathcal{F})$$

of local traces is equal to the alternating sum of traces of the *global* geometric Frobenius automorphism of U acting on the compactly-supported étale

⁶ Whether or not they have finite image.

cohomology groups $H_c^i(\overline{U}, \mathcal{F})$ of the sheaf (see, e.g., [51], [29], [97, VI.13]). Hence we have

$$W(\pi, \tau) = \frac{1}{\sqrt{|\hat{\Gamma}_m^\pi| |\hat{\Gamma}_n^\tau|}} \sum_{i=0}^{2d} (-1)^i \operatorname{Tr}(\operatorname{Fr} \mid H_c^i(\overline{U}, \mathcal{W}(\pi, \tau))),$$

where Fr denotes the global geometric Frobenius; recall that d is the dimension of U .

Since the representation corresponding to $\mathcal{W}(\pi, \tau)$ factors through a finite group, any eigenvalue of any $\operatorname{Fr}_{u, q^n}$ for $u \in U(\mathbf{F}_{q^n})$ is a root of unity, so this sheaf is *pointwise pure* of weight 0. Therefore, by Deligne's extraordinary generalization of the Riemann Hypothesis over finite fields (see [28, p. 138], [29, Theorem 1.17]), all the eigenvalues of the geometric Frobenius automorphism Fr acting on $H_c^i(\overline{U}, \mathcal{W}(\pi, \tau))$ are algebraic integers, and all the Galois-conjugates of a given eigenvalue are of absolute value $q^{w/2}$ for some integer $w \leq i/2$ (called the weight, which is independent of the conjugate, for a given eigenvalue, but may depend on the eigenvalue itself).

In particular, each term in the alternating sum is an algebraic number, and we may see this as a formula valid in \mathbf{C} (after fixing some embedding of the algebraic closure of \mathbf{Q} inside \mathbf{C}), so that speaking of the modulus of the terms makes sense. Isolating the contribution of the topmost cohomology group $H_c^{2d}(\overline{U}, \mathcal{W}(\pi, \tau))$, this leads to

$$W(\pi, \tau) = \frac{1}{\sqrt{|\hat{\Gamma}_m^\pi| |\hat{\Gamma}_n^\tau|}} \operatorname{Tr}(\operatorname{Fr} \mid H_c^{2d}(\overline{U}, \mathcal{W}(\pi, \tau))) + O\left(\sigma'_c(\overline{U}, \mathcal{W}(\pi, \tau)) q^{d-1/2}\right),$$

where the implied constant is ≤ 1 and where

$$\sigma'_c(\overline{U}, \mathcal{W}(\pi, \tau)) = \sum_{i=0}^{2d-1} \dim H_c^i(\overline{U}, \mathcal{W}(\pi, \tau)).$$

For the 'main term', we use the coinvariant⁷ formula (see, e.g., [29, Sommes trigonométriques, Remarque 1.18d])

$$H_c^{2d}(\overline{U}, \mathcal{W}(\pi, \tau)) = V_{\pi_1(\overline{U}, \bar{\eta})}(-d)$$

where $V = \mathcal{W}(\pi, \tau)_{\bar{\eta}}$ is the space (i.e., the fiber over $\bar{\eta}$) on which the representation which 'is' the sheaf acts, and $(-d)$ denotes a Tate twist (which means that the eigenvalues of Fr on the cohomology group are q^d times the eigenvalues of Fr as it acts naturally on the coinvariant space). By the linear disjointness

⁷ Recall that for a vector space V over a field, on which a group G acts, the coinvariant space V_G is the largest quotient on which G acts trivially, in other words, $V_G = V / \langle (g - 1)v \rangle$.

assumption, when we factor $[\pi, \bar{\tau}]$ restricted to the geometric fundamental group as follows

$$\pi_1(\bar{U}, \bar{\eta}) \xrightarrow{\rho_{[m,n]}} G_{[m,n]}^g \xrightarrow{[\pi, \bar{\tau}]} GL((\dim \pi)(\dim \tau), \bar{\mathbf{Q}}_\ell),$$

the first map is *surjective*. Hence we have

$$V_{\pi_1(\bar{U}, \bar{\eta})}(-d) = W_{G_{[m,n]}^g}(-d)$$

with W denoting the space of $[\pi, \bar{\tau}]$. As we are dealing with linear representations of finite groups in characteristic 0, this last coinvariant space is isomorphic to the space of invariants under $G_{[m,n]}^g$, and in particular its dimension is the multiplicity of the trivial representation in $[\pi, \bar{\tau}]$ restricted to $G_{[m,n]}^g$. By Lemma 3.4, we have therefore

$$H_c^{2d}(\bar{U}, \mathcal{W}(\pi, \tau)) = 0$$

if $(m, \pi) \neq (n, \tau)$.

If $(m, \pi) = (n, \tau)$, on the other hand, the same lemma states that the dimension of $H_c^{2d}(\bar{U}, \mathcal{W}(\pi, \pi))$ is equal to $|\hat{\Gamma}_m^\pi|$. Now we claim that the global Frobenius acts *trivially* on the invariant space, and so by multiplication by q^d on the cohomology group because of the Tate twist. Indeed, although the Frobenius is *not* in $\pi_1(\bar{U}, \bar{\eta})$, it acts with finite order as the topological generator -1 of the quotient $G/G^g \simeq \hat{\mathbf{Z}}$, and if we decompose $[\pi, \bar{\pi}]^{G_m^g} = (\pi \otimes \bar{\pi})^{G_m^g}$ as a direct sum of characters of G_m/G_m^g , which is a finite cyclic group, we obtain as in Lemma 3.4 that this invariant space is isomorphic to the direct sum of the characters $\psi \in \hat{\Gamma}_m^\pi$, on each of which Fr acts by multiplication by $\psi(-1)$. Now the relation $\pi \simeq \pi \otimes \psi$ for $\psi \in \hat{\Gamma}_m^\pi$ provides for *every* $x \in Y_m$ (hence with $\psi(x) = \psi(d(x)) = \psi(-1)$) the relation

$$\text{Tr } \pi(x) = \psi(-1) \text{Tr } \pi(x),$$

so that $\psi(-1) = 1$, as otherwise $\text{Tr } \pi$ would vanish identically on Y_m , which we excluded in the definition of Π_m^* .

This evaluation of the first term gives now

$$W(\pi, \tau) = \delta((m, \pi), (n, \tau))q^d + O\left(\sigma'_c(\bar{U}, \mathcal{W}(\pi, \tau))q^{d-1/2}\right),$$

with an absolute implied constant.

To conclude, we need a bound for $\sigma'_c(\bar{U}, \mathcal{W}(\pi, \tau))$, uniform in terms of m, n, π and τ . This is where the distinction between the two cases of the proposition occur. Those bounds were proved in [80, Proposition 4.1; Proposition 4.7], and we will explain them for completeness in Section 8.4 below (with small changes). In particular, quoting Proposition 8.9 suffices to conclude the proof. \square

Readers who are primarily interested in the applications of the sieve for Frobenius are invited to skip the next section, which is the most heavily dependent on fairly advanced techniques of algebraic geometry (possibly after checking that Proposition 8.9 does indeed apply).

8.4 Estimates for sums of Betti numbers

The dimensions of cohomology groups of a sheaf are called its *Betti numbers*. In this section, we are interested in the sum of the Betti numbers $\sigma'_c(\overline{U}, \mathcal{W}(\pi, \tau))$ which appeared in the previous section. We state and explain the proof of slightly more general estimates for such sums, since those (or their proof) may be of independent interest (for instance, the argument in Case (1) is used in [1] and [2]). For more details and references, see Section 4 in [80].

For a $\overline{\mathbf{Q}}_c$ -adic sheaf \mathcal{F} on U , we write

$$\sigma_c(\overline{U}, \mathcal{F}) = \sum_{i=0}^{2d} \dim H_c^i(\overline{U}, \mathcal{F}).$$

Proposition 8.9 *Let q be a power of a prime p , let U/\mathbf{F}_q be a smooth affine geometrically irreducible algebraic variety of dimension $d \geq 1$.*

- (1) *Let $\rho : \pi_1(U, \overline{\eta}) \rightarrow G$ be a continuous surjective homomorphism with G finite of order prime to p and let $\pi : G \rightarrow GL(r, \overline{\mathbf{Q}}_\ell)$ be a representation of G for some $\ell \neq p$. Denote by $\pi(\rho)$ the lisse sheaf on U associated to $\pi \circ \rho$. There exists a constant $C(\overline{U})$ depending only on \overline{U} such that*

$$\sigma_c(\overline{U}, \pi(\rho)) \leq C(\overline{U})|G|(\dim \pi). \quad (8.7)$$

In fact, if $d = 1$ we can take $C(\overline{U}) = \sigma_c(\overline{U}, \mathbf{Q}_\ell)$, and if $d \geq 2$ and U is embedded in the affine space of dimension N using r equations of degree $\leq \delta$, we can take

$$C(\overline{U}) = C(N, r, \delta) \leq 12N2^r(3 + r\delta)^{N+1}. \quad (8.8)$$

- (2) *Let $d = 1$ and assume that $U = C - S$ is the complement of a non-empty finite set of points S in a smooth projective curve C/\mathbf{F}_q . Let $\rho : \pi_1(U, \overline{\eta}) \rightarrow G$ be a continuous surjective map with*

$$G = \prod_{1 \leq i \leq k} G_i, \quad \rho = \rho_1 \otimes \cdots \otimes \rho_k,$$

where G_i , $1 \leq i \leq k$, is a subgroup of $GL(r, k_i)$ with k_i a finite field of characteristic $\ell_i \neq p$, and ρ_i is a map from $\pi_1(U, \overline{\eta})$ to G_i such that (ρ_i) is

part of a compatible system. Denote by $\pi(\rho)$ the lisse sheaf on U associated to $\pi \circ \rho$. Then we have

$$\sigma'_c(\overline{U}, \pi(\rho)) \leq C(\overline{U}, (\rho_i))(\dim \pi),$$

for some constant $C(\overline{U}, (\rho_i))$ depending only on \overline{U} and the compatible system. In fact, we can take

$$C = 1 - \chi_c(\overline{U}, \mathbf{Q}_\ell) + |S|w \quad (8.9)$$

where $w \geq 1$ is the sum of the Swan conductors of all \mathcal{F}_i at the points in S , which is independent of i .

Note that in Case (2), the statement in [80] is slightly different, giving $C = 1 - \chi_c(\overline{U}, \mathbf{Q}_\ell) + kw$ instead of (8.9). This may be more useful if the size of S is large (if applied to a sequence of curves with $|S| \rightarrow +\infty$), but is worse when k gets large, as is the case when using a sieve support with m having possibly many prime factors. Precisely, this leads to a loss of a power of $\log \log L$ in some applications as in Section 8.6; some readers may consider this loss to be well within reason . . .

Proof In Case (1), there are three basic tools, which exploit deeply the powerful formalism of étale cohomology:

- Consider the étale covering $\overline{V} \rightarrow \overline{U}$ corresponding to the kernel of $\pi \circ \rho$ (restricted to the geometric fundamental group). A standard property of étale cohomology is that we have

$$\dim H_c^i(\overline{U}, \pi(\rho)) \leq \dim H_c^i(\overline{V}, \overline{\mathbf{Q}}_\ell^{\dim \pi}) = (\dim \pi) \dim H_c^i(\overline{V}, \overline{\mathbf{Q}}_\ell),$$

for all i , and hence

$$\sigma_c(\overline{U}, \pi(\rho)) \leq (\dim \pi) \sigma_c(\overline{V}).$$

Moreover the covering $\overline{V} \rightarrow \overline{U}$ is tamely ramified by assumption on the image of ρ , and therefore it is sufficient to show that given a tamely ramified Galois covering $\overline{V} \rightarrow \overline{U}$ with Galois group G , where U/\mathbf{F}_q is smooth affine and geometrically irreducible of dimension $d \geq 1$, we have

$$\sigma_c(\overline{V}) \leq C(\overline{U})|G|.$$

(Note that tame ramification is a necessary condition for the existence of such a general bound; there are counterexamples otherwise, see [80, Remark 4.8]).

- An adaptation of a method of Katz [71] allows an argument by induction on d , by relating $\sigma_c(\overline{V})$ to the Euler–Poincaré characteristic

$$\chi_c(\overline{V}) = \sum_{i=0}^{2d} (-1)^i \dim H_c^i(\overline{V}, \overline{\mathbf{Q}}_\ell).$$

- A result due to Deligne and Lusztig (see [30, 3.12], or the explanation in [64, 2.6, Corollary 2.8]) shows that the Euler–Poincaré characteristic satisfies

$$\chi_c(\overline{V}) = |G| \chi_c(\overline{U})$$

for a tamely ramified étale covering $\overline{V} \xrightarrow{\varphi} \overline{U}$ with Galois group G , so this reduces in the induction to a problem over the base \overline{U} ; the precise induction step follows ideas of Katz and is based on a Bertini-type theorem which states that one can find a hyperplane section $U \cap H$ of U such that the inverse image W of V over $U \cap H$ is still a (tamely ramified) connected Galois covering with Galois group G , and moreover the cohomology groups of V are sufficiently controlled by those of W .

Combining these ingredients, the result follows, and we refer to [80, Section 4] for details.

In Case (2), which corresponds to Proposition 4.1 of [80], what is needed is the analysis of the ramification of sheaves on open curves, which is described precisely in [72, Chapter 1].

Since the curve \overline{U} is affine and smooth, there are no compactly supported sections of the sheaf $\pi(\rho)$, so that $H_c^0(\overline{U}, \pi(\rho)) = 0$, and hence

$$\sigma'_c(\overline{U}, \pi(\rho)) = \dim H_c^1(\overline{U}, \pi(\rho)) = \dim H_c^2(\overline{U}, \pi(\rho)) - \chi_c(\overline{U}, \pi(\rho)).$$

In order to find an upper bound for $-\chi_c(\overline{U}, \pi(\rho))$, we apply the Euler–Poincaré formula of Grothendieck–Ogg–Shafarevitch (see, e.g., [72, 2.3.1, 2.3.3]), which gives

$$\chi_c(\overline{U}, \pi(\rho)) = (\dim \pi) \chi_c(\overline{U}) - \sum_{x \in S} \text{Swan}_x(\pi(\rho)),$$

where S is the set of ‘points at infinity’, and $\text{Swan}_x(\pi(\rho))$ is the Swan conductor of the sheaf $\pi(\rho)$ at x , a certain non-negative number which is defined using the action of $\pi(\rho)$ on the higher-ramification groups of the inertia group at x . (In particular, $\text{Swan}_x(\pi(\rho)) = 0$ for all x if and only if $\pi(\rho)$ is tamely ramified, in which case the formula above is a special case of the result of Deligne–Lusztig quoted above, and what follows is much easier).

From the description of ρ as a tensor product of representations on k_i -vector spaces, fairly simple properties of the Swan conductor imply that

$$\mathrm{Swan}_x(\pi(\rho)) \leq \mathrm{Swan}_x(\rho) \leq \max_{1 \leq i \leq k} \mathrm{Swan}_x(\rho_i),$$

and therefore

$$-\chi_c(\overline{U}, \pi(\rho)) \leq (\dim \pi)(-\chi_c(\overline{U})) + \sum_{x \in S} \max_{1 \leq i \leq k} \mathrm{Swan}_x(\rho_i). \quad (8.10)$$

Now we exploit the fact that there is a compatible system $(\tilde{\rho}_i)$ which yields (ρ_i) by reduction. Another standard property of the Swan conductor is that

$$\mathrm{Swan}_x(\rho_i) = \mathrm{Swan}_x(\tilde{\rho}_i).$$

Moreover, since the Euler–Poincaré characteristic of $\tilde{\rho}_i$ is independent of i (being minus the degree of the L -function associated to $\tilde{\rho}_i$), the quantity

$$w = \sum_{x \in S} \mathrm{Swan}_x(\tilde{\rho}_i) = r \chi_c(\overline{U}) - \chi_c(\overline{U}, \tilde{\rho}_i)$$

is independent of i . Now notice that for any i and fixed $x_0 \in S$, we have

$$\mathrm{Swan}_{x_0}(\tilde{\rho}_i) \leq \sum_x \mathrm{Swan}_x(\tilde{\rho}_i) = w,$$

so that

$$\max_{1 \leq i \leq k} \mathrm{Swan}_{x_0}(\tilde{\rho}_i) \leq w.$$

Inserting this bound in (8.10) we obtain

$$-\chi_c(\overline{U}, \pi(\rho)) \leq (\dim \pi)(-\chi_c(\overline{U})) + |S|w$$

and it suffices to add the contribution of $H_c^2(\overline{U}, \pi(\rho))$, which is at most $\dim \pi$ by the coinvariant formula already used earlier, to obtain the stated bound for $\sigma'_c(\overline{U}, \pi(\rho))$. \square

8.5 Bounds for the large sieve constants

In order to apply the bounds for the exponential sums given by Proposition 8.8 to the estimation of the large sieve constants for the sieve for Frobenius, we see that we will need upper bounds for the quantities

$$\max_{m, \pi} \left\{ q^d + C q^{d-1/2} (\dim \pi) \sum_{n \leq L} |G_{[m, n]}| \sum_{\tau \in \Pi_n^*} (\dim \tau) \right\} \quad (8.11)$$

in the first case and

$$\max_{m, \pi} \left\{ q^d + Cq^{d-1/2}(\dim \pi) \sum_{n \leq L}^b \sum_{\tau \in \Pi_n^*} (\dim \tau) \right\} \quad (8.12)$$

in the second case.

For this purpose, we make the following assumptions: for all $\ell \in \Lambda$, and $\pi \in \Pi_\ell^*$, we have

$$|G_\ell| \leq (\ell + 1)^s, \quad \dim \pi \leq (\ell + 1)^v, \quad \sum_{\pi \in \Pi_\ell^*} (\dim \pi) \leq (\ell + 1)^t, \quad (8.13)$$

where s , t and v are non-negative integers. In the notation of Chapter 5, the second and third are implied by

$$A_\infty(G_\ell) \leq (\ell + 1)^v, \quad A_1(G_\ell) \leq (\ell + 1)^t$$

respectively; indeed, the results of Chapter 5 were motivated by the desire to have optimal bounds of this type for certain specific finite groups of Lie type which are encountered in applications.

Here are important special cases that follow from Example 5.8 in Chapter 5 (and from the character table of $GL(2, \mathbf{F}_\ell)$, which we have included for convenience at the end of Appendix B):

- If G_ℓ is a subgroup of $GL(r, \mathbf{F}_\ell)$, we can take $s = r^2$, $v = r(r - 1)/2$, $t = r(r + 1)/2$.
- If G_ℓ is a subgroup of symplectic similitudes for some non-degenerate alternating form of rank $2g$, we can take $s = g(2g + 1) + 1$, $v = (s - (g + 1))/2 = g^2$, $t = g^2 + g + 1$.
- In particular, if $G_\ell \subset GL(2, \mathbf{F}_\ell)$ and $G_\ell^s = SL(2, \mathbf{F}_\ell)$, we have

$$|G_\ell| \leq \ell^4, \quad \max(\dim \pi) = \ell + 1, \quad \sum_{\pi \in \Pi_\ell^*} (\dim \pi) \leq (\ell + 1)^3. \quad (8.14)$$

Recall the definition of the arithmetic function $\psi(m)$, namely

$$\psi(m) = \prod_{\ell|m} (\ell + 1)$$

for $m \geq 1$. From (8.13), we deduce by multiplicativity that we have

$$|G_m| \leq \psi(m)^s, \quad \dim \pi \leq \psi(m)^v, \quad \sum_{\pi \in \Pi_m^*} (\dim \pi) \leq \psi(m)^t, \quad (8.15)$$

for all squarefree integers $m \geq 1$.

We wish to sieve with the prime sieve support $\mathcal{L}^* = \{\ell \in \Lambda \mid \ell \leq L\}$ for some L . The first idea for the sieve support itself is to use the traditional one, say \mathcal{L}_1 , namely the set of squarefree integers $m \leq L$ divisible only by primes in

Λ . However, since we have $\psi(m) \ll m \log \log m$, and this upper bound is sharp (if m has many small prime factors), the use of \mathcal{L}_1 leads to a loss of a power of $\log \log L$ in the second term in the estimation of (8.11) and (8.12). As described by D. Zywna (in his preprint ‘The large sieve and Galois representations’, 2007), this can be recovered using the trick of sieving using only squarefree integers m which are free of small prime factors, in the sense that $\psi(m) \leq L+1$ instead of $m \leq L$ (which for primes ℓ remains equivalent with $\ell \leq L$). This means we use the sieve support

$$\mathcal{L} = \{m \in S(\Lambda) \mid m \text{ is squarefree and } \psi(m) \leq L+1\}.$$

We quote both types of sieves:

Corollary 8.10 *With the above data and notation, in particular under the assumption of linear disjointness of the system (ρ_ℓ) for $\ell \in \Lambda$, let $\Omega_\ell \subset G_\ell$, for all primes $\ell \in \Lambda$, be a conjugacy-invariant subset of G_ℓ such that $d(\Omega_\ell) = -1$. Then there exists a constant $C \geq 0$ such that we have both*

$$|\{u \in U(\mathbf{F}_q) \mid \rho_\ell(F_u) \notin \Omega_\ell \text{ for } \ell \leq L\}| \leq (q^d + Cq^{d-1/2}(L+1)^A)H^{-1} \quad (8.16)$$

and

$$\begin{aligned} |\{u \in U(\mathbf{F}_q) \mid \rho_\ell(F_u) \notin \Omega_\ell \text{ for } \ell \leq L\}| \\ \leq (q^d + Cq^{d-1/2}L^A(\log \log L)^B)K^{-1}, \end{aligned} \quad (8.17)$$

where

$$H = \sum_{\substack{\psi(m) \leq L+1 \\ \ell \mid m}}^b \prod \frac{|\Omega_\ell|}{|G_\ell^g| - |\Omega_\ell|}, \quad K = \sum_{m \leq L}^b \prod_{\ell \mid m} \frac{|\Omega_\ell|}{|G_\ell^g| - |\Omega_\ell|},$$

and

- (i) If $p \nmid |G_\ell|$ for all $\ell \in \Lambda$, we can take $A = v+2s+t+1$ and $B = v+s$, and the constant C depends only on \overline{U} and on s , t and v in the case of (8.17).
- (ii) If $d = 1$ and the system (ρ_ℓ) arises by reduction of a compatible system of \mathbf{Z}_ℓ -adic sheaves on U , then we can take $A = t+v+1$ and $B = v$, and the constant C depends only on the Euler–Poincaré characteristic of \overline{U} , the compactly-supported Euler–Poincaré characteristic of the compatible system (\tilde{W}_ℓ) on \overline{U} , and on s , t , v in the case of (8.17).

Proof From Proposition 2.9, we must estimate

$$\Delta = \max_{m, \pi} \sum_n^b \sum_{\tau \in \Pi_n^*} |W(\pi, \tau)|,$$

where m and n run over \mathcal{L} (or \mathcal{L}_1). By Proposition 8.8, this is bounded by the quantities (8.11) and (8.12). Using (8.15), the result is now straightforward, using (in the case of (8.17)) the simple estimate

$$\sum_{n \leq L}^b \psi(n)^A \ll L^{A+1}$$

for $L \geq 1$, $A \geq 0$, the implied constant depending only on A (for readers unfamiliar with this type of analytic number theory estimate, we sketch the proof in Proposition G.3 of Appendix G). \square

Remark 8.11 In [80], we used assumptions on the size of the monodromy groups and the dimensions of their representations which were different from (8.13), precisely we assumed

$$|G_\ell| \leq c_1 \ell^s, \quad |G_\ell^\#| \leq c_2 \ell^t$$

for some constants $c_1, c_2 \geq 0$ and $s, t \geq 0$. The crucial feature of the current assumptions is that $A_1(G_\ell)$ and $A_\infty(G_\ell)$ are bounded by *monic* polynomials in ℓ . Having polynomials with constant terms > 1 would mean, after multiplicativity is applied, that $A_\infty(G_m)$ and $A_1(G_m)$ would be bounded by polynomials times a divisor function: for instance, $A_1(G_\ell) \leq c_3 \ell^v$ with $c_3 \geq 1$ implies

$$A_1(G_m) \leq c_3^{\omega(m)} m^v,$$

and on average over m , this would mean a loss of a power of logarithm since we have

$$\sum_{m \leq L}^b c_3^{\omega(m)} m^v \asymp L^{v+1} (\log L)^{c_3-1}, \quad \text{as } L \rightarrow +\infty$$

(see Appendix G). When applying the sieve in a genuinely large sieve situation (as above with irreducibility of zeta functions of curves), the effect of losing a power of $\log L$ in the numerator of the sieve bound overwhelms the corresponding saving coming from the use of squarefree numbers in the denominator (typically, a power of $\log L$ with exponent < 1 ; compare the order of magnitude arising from hypothetical bounds

$$\frac{q + C\sqrt{q}L(\log L)}{\sum_{m \leq L}^b 2^{-\omega(m)}}$$

and

$$\frac{q + C\sqrt{q}L}{\sum_{\ell \leq L} \frac{1}{2}}$$

where the former yields only $\sqrt{q}(\log q)^{3/2}$ when $L = \sqrt{q}(\log q)^{-1}$, while the latter gives $\sqrt{q}(\log q)$ with $L = \sqrt{q}$. Hence, bounds such as (8.13) are necessary to benefit from the enlarged sieve support.

On the other hand, in ‘small sieve’ settings, the gain from the use of squarefree numbers is much more important (since the size of H typically grows from a multiple of $\log \log L$ to a power of $\log L$), and this may well be sufficient justification for using simpler but weaker polynomial bounds in such cases.

8.6 Application to Chavdarov's problem

Using the sieve for Frobenius, we can give a strong answer to the problem solved qualitatively by Chavdarov, in particular in the case of the family of curves described in Example 8.6. The idea is to apply the same general ideas as Gallagher's Theorem, and this means that we need to control the distribution of the reductions modulo primes of the polynomials that we want to study. This solution depends on the following crucial facts which explain precisely why the problem is amenable to the sieve we have described (in particular, looking at ‘algebraic’ families of zeta functions of curves is essential):⁸

- If C/\mathbf{F}_q is a smooth projective geometrically connected algebraic curve over a finite field, the numerator P of the zeta function of C described in Theorem 8.1 is given by

$$P(T) = \det(1 - T \operatorname{Fr} \mid H_c^1(C \times \overline{\mathbf{F}}_q, \mathbf{Z}_\ell)) \quad (8.18)$$

where ℓ is an arbitrary prime different from p .

- Moreover, if $\mathcal{C} \xrightarrow{\pi} U$ is an algebraic family of smooth geometrically connected projective curves of genus g , as described in Section 8.1, the variation of the polynomial is captured by a family of lisse ℓ -adic sheaves on U , technically known as the first higher direct images with compact support of the trivial sheaf on U and denoted $\mathcal{F}_\ell = R^1 \pi_* \mathbf{Z}_\ell$: for every $n \geq 1$ and $t \in U(\mathbf{F}_{q^n})$, we have

$$P_t = \det(1 - T \operatorname{Fr}_{t, q^n} \mid \mathcal{F}_\ell),$$

or in other words, there exists a free \mathbf{Z}_ℓ -module of rank $2g$ with a continuous action $\tilde{\rho}_\ell$ of $\pi_1(U, \bar{\eta})$ for which

$$P_t = \det(1 - T \tilde{\rho}(\operatorname{Fr}_{t, q^n})).$$

⁸ This is not meant to imply that other ways of putting together algebraic curves might not lead to similar properties, simply that the sieve for Frobenius would probably not be the right tool for the job. See for instance [84, Section 4] for a study of *isogeny classes* of abelian varieties, instead of an algebraic family, where the classical large sieve is used.

Since the numerator of the zeta function of a curve is a polynomial with integer coefficients defined without reference to any auxiliary prime ℓ , note that this means in particular that the family $(\tilde{\rho}_\ell)$ is a compatible system.

Moreover, the functional equation (8.3) may be deduced from the existence of a non-degenerate alternating pairing $\langle \cdot, \cdot \rangle$ on $H_c^1(C \times \overline{\mathbf{F}}_q, \mathbf{Z}_\ell)$, coming from Poincaré duality,⁹ such that the global Frobenius acts as a *symplectic similitude* with multiplier q , i.e., for all v, w , we have

$$\langle \text{Fr}(v), \text{Fr}(w) \rangle = q \langle v, w \rangle;$$

in turn, in the case of a family of curves, this means that the representation $\tilde{\rho}_\ell$ takes value in the group $CSp(\langle \cdot, \cdot \rangle) \simeq CSp(2g, \mathbf{Z}_\ell)$ of symplectic similitudes for this pairing, and that for any $n \geq 1$ and $t \in U(\mathbf{F}_{q^n})$, the image $\tilde{\rho}_\ell(\text{Fr}_{t,q^n})$ is a symplectic similitude with multiplier q^n . We have in fact a commutative diagram with exact rows

$$\begin{array}{ccccccccc} 1 & \longrightarrow & \pi_1(\overline{U}, \bar{\eta}) & \longrightarrow & \pi_1(U, \bar{\eta}) & \xrightarrow{d} & \hat{\mathbf{Z}} & \longrightarrow & 1 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 1 & \longrightarrow & Sp(2g, \mathbf{Z}_\ell) & \longrightarrow & CSp(2g, \mathbf{Z}_\ell) & \xrightarrow{m} & \mathbf{Z}_\ell^\times & \longrightarrow & 1, \end{array}$$

(but the vertical arrows are not always surjective).

Note that this interpretation of the zeta function combined with Proposition E.1 in Appendix E explains again why the Galois group of the splitting field of P_t can be seen as a subgroup of W_{2g} .

Furthermore, we now see that we can control the reduction of the polynomials P_t modulo a prime $\ell \neq p$ by looking at the maps induced from $\tilde{\rho}_\ell$ by reduction modulo ℓ : let

$$\rho_\ell : \pi_1(U, \bar{\eta}) \rightarrow CSp(2g, \mathbf{F}_\ell) \subset GL(2g, \mathbf{F}_\ell),$$

then

$$P_t \pmod{\ell} = \det(1 - T\rho_\ell(\text{Fr}_{t,q^n})) \in \mathbf{F}_\ell[T], \quad (8.19)$$

for all $\ell \neq p$. This family (ρ_ℓ) forms a family of group homomorphisms of the type described in the general setting of the sieve for Frobenius in Section 8.2 (in fact, a compatible system), and provides the required link between our concrete diophantine problem and the algebraic geometry discussed previously . . .

At this point, for a given family of curves, the issue is clear: to apply the sieve for Frobenius, it is necessary to determine, as precisely as possible, the image $G_\ell = \rho_\ell(\pi_1(U, \bar{\eta}))$ of ρ_ℓ , and more particularly to see whether (or not)

⁹ This is the analogue of the intersection pairing on closed surfaces of genus g that occurred in the previous chapter.

the linear disjointness condition, i.e., the surjectivity of the maps (8.6), holds. (Clearly *some* condition is needed, because we can always take a 'trivial' family $\mathcal{C} = C \times U$ with projection onto U , for some fixed curve C , and if the numerator of the zeta function of the latter is not irreducible, none of those of curves in this family will be.) Intuitively, linear disjointness can be expected to hold if the family 'varies a lot', and in particular it holds for purely group-theoretical reasons if the geometric monodromy groups $G_\ell^g = \rho_\ell(\pi_1(\bar{U}, \bar{\eta}))$ are as large as possible, namely if

$$G_\ell^g = Sp(2g, \mathbf{F}_\ell) \quad (8.20)$$

for all $\ell \neq p$ (see the commutative diagram above to check that the image of the geometric fundamental group by ρ_ℓ is inside the kernel $Sp(2g, \mathbf{F}_\ell)$ of the multiplier map; also it would be sufficient that this holds for almost all ℓ , in some sense, since a few exceptions will not matter in applying the sieve). This is the content of the following lemma:

Lemma 8.12 *Let m be a squarefree integer, $g \geq 1$ an integer, H a subgroup of the product*

$$G = \prod_{\ell|m} Sp(2g, \mathbf{F}_\ell)$$

which maps onto each factor $Sp(2g, \mathbf{F}_\ell)$ for $\ell \mid m$. Then we have $H = G$.

This is a consequence of a variant of Goursat's lemma, and is proved for instance by Chavdarov in [22, Proposition 5.1].

It may seem that the maximality condition (8.20) is very restrictive, and maybe impossible to verify;¹⁰ however, although it is indeed a delicate issue, it turns out that there are quite a few cases where the condition holds, and can be checked.

We will comment further on this below, but for the moment we indicate one particular case where G_ℓ^g is well understood, already used by [22] and corresponding to Example 8.6. This is a theorem of J.-K. Yu ('Toward a proof of the Cohen–Lenstra conjecture in the function field case', preprint, 1996):

Proposition 8.13 *Let q be a power of an odd prime number, let $g \geq 1$ be an integer, and let $f \in \mathbf{F}_q[X]$ be a squarefree monic polynomial of degree $2g$. Let U/\mathbf{F}_q be the open subset of the affine line where f does not vanish, and let $\mathcal{C} \xrightarrow{\pi} U$ be the family of smooth projective geometrically connected*

¹⁰ Those who believe in the field with one element will indeed see that the desired outcome of all this, that P_t has 'generically' maximal Galois group W_{2g} , looks suspiciously like the *same* statement of maximality for this mythic beast instead of $\mathbf{F}_\ell \dots$

hyperelliptic curves of genus g given by the smooth projective models of the affine curves with equations

$$y^2 = f(x)(x - t).$$

Then the map $\rho_\ell : \pi_1(\overline{U}, \bar{\eta}) \rightarrow Sp(2g, \mathbf{F}_\ell)$ is onto for all odd primes ℓ .

Remark 8.14 The case $\ell = 2$ must be excluded, because roots of f provide rational 2-torsion points of the Jacobian of \mathcal{C} which are invariant under G_2^g , so that this group may fail to be maximal.

Yu's proof (which works by reduction to characteristic 0) is still unpublished, but there are two recent proofs, one by C. Hall [53] (who uses methods related to those developed by Katz to compute the rational monodromy,¹¹ and gives fairly general criteria to show that the finite geometric monodromy groups of a family of sheaves are 'large', the proposition being only a very special case of his work; see also [82] for a write-up of Hall's theorem in this special case), and the other by J. Achter and R. Pries [3] with a more algebro-geometric flavour (moduli spaces and study of degenerations of curves to argue by induction on g).

We now have all ingredients available to prove a version of the quantitative solution to Problem 8.5. We do this here for the curves of Example 8.6; in the next section, we will comment on more general versions.

Theorem 8.15 *Let \mathbf{F}_q be a finite field of characteristic $p \neq 2$, let $f \in \mathbf{F}_q[X]$ be a squarefree monic polynomial of degree $2g$, $g \geq 1$. For $t \in \mathbf{F}_q$ which is not a zero of f , let $P_t \in \mathbf{Z}[T]$ be the numerator of the zeta function of the smooth projective model of the hyperelliptic curve*

$$C_t : y^2 = f(x)(x - t), \quad (8.21)$$

and let K_t be the splitting field of P_t over \mathbf{Q} , which has degree $[K_t : \mathbf{Q}] \leq |W_{2g}| = 2^s g!$. Then we have

$$|\{t \in \mathbf{F}_q \mid f(t) \neq 0 \text{ and } [K_t : \mathbf{Q}] < 2^s g!\}| \ll q^{1-\gamma} (\log q)^{1-\delta}$$

where $\gamma = (4g^2 + 2g + 4)^{-1}$ and $\delta > 0$, with $\delta \sim 1/(4g)$ as $g \rightarrow +\infty$. The implied constant depends only on g , and in particular the estimate is valid with q replaced by q^n for all $n \geq 1$.

¹¹ That is, the Zariski closure of the image of the geometric fundamental group by $\tilde{\rho}_\ell$, see the discussion in Section 8.7.

Proof Let U be the open set of the affine line over \mathbf{F}_q where f does not vanish. According to the previous discussion, and in parallel with the argument in Theorem 4.2, we can apply the sieve for Frobenius to the system (ρ_ℓ) arising by reduction from the compatible system $(\tilde{\rho}_\ell)$, defined for any odd prime $\ell \neq p$, which is linearly disjoint by Yu's theorem. From Corollary 8.10, in Case (ii) with the first bound, applied with

$$\mathcal{L}^* = \{\ell \mid 3 \leq \ell \leq L, \ell \neq p\}, \quad \mathcal{L} = \{m \in S(\Lambda) \mid \psi(m) \leq L+1\},$$

we have indeed

$$|\{u \in U(\mathbf{F}_q) \mid \rho_\ell(F_u) \notin \Omega_\ell \text{ for } \ell \in \mathcal{L}\}| \leq \left(q + (2g-1)q^{1/2}(L+1)^{2g^2+g+2}\right)H^{-1}$$

for any choice of subsets $\Omega_\ell \subset CSp(2g, \mathbf{F}_\ell)$ such that the multiplier of Ω_ℓ is always q , with

$$H = \sum_{\substack{\psi(m) \leq L+1 \\ (m, 2p)=1}}^b \prod_{\ell \mid m} \frac{|\Omega_\ell|}{|Sp(2g, \mathbf{F}_\ell)| - |\Omega_\ell|};$$

in applying the upper-bound, we have taken $A = g^2 + g + 1 + g^2 + 1 = 2g^2 + g + 2$ by Proposition 8.13 and Section 8.5, and $C = 2g - 1$ by looking at Proposition 8.9, (2): it is known that the sheaves $R^1\pi_1\mathbf{Z}_\ell$ on U are tame (this is shown for instance in [77, Lemma 10.1.12]), so that $w = 0$, while $\chi_c(\overline{U}, \mathbf{Q}_\ell) = 2 - 2g$ by additivity of the Euler–Poincaré characteristic ($\chi_c(\overline{U}) + \chi_c(\mathbf{P}^1 - \overline{U}) = \chi_c(\mathbf{P}^1) = 2$, and $\chi_c(\mathbf{P}^1 - \overline{U}) = 2g$ since this is a zero-dimensional variety).

Now in order to apply the sieve to Chavdarov's problem, we appeal to the principle (already present in Gallagher's Theorem) that the factorization of a polynomial $f \in \mathbf{F}_\ell[T]$ in irreducible factors gives indication on which conjugacy classes of permutations the Galois group of the splitting field of f contains. Specifically, according to the result of Exercise 8.1, for $t \in U(\mathbf{F}_q)$, if the splitting field of the numerator P_t of the zeta function of C_t is not maximal, then it follows that, for some $i = 1, 2, 3, 4$ (or $i = 1$ or 2 if $g = 1$), we have

$$\rho_\ell(\text{Fr}_{t,q}) \notin \Omega_{i,\ell}$$

for all primes $\ell \nmid 2p$, where:

- $\Omega_{1,\ell}$ is the set of $g \in CSp(2g, \mathbf{F}_\ell)$ with multiplier $m(g) = q$ such that the reversed characteristic polynomial $\det(1 - Tg) \in \mathbf{F}_\ell[T]$ is irreducible.
- $\Omega_{2,\ell}$ is the set of $g \in CSp(2g, \mathbf{F}_\ell)$ with $m(g) = q$ such that $\det(1 - Tg)$ factors as a product of an irreducible quadratic polynomial and a product of irreducible polynomials of odd degrees.

- If $g \geq 2$, $\Omega_{3,\ell}$ is the set of $g \in CSp(2g, \mathbf{F}_\ell)$ with $m(g) = q$ such that if we factor¹²

$$\det(1 - Tg) = T^g h(qT + T^{-1}) \quad (8.22)$$

with $h \in \mathbf{F}_\ell[T]$ a monic polynomial of degree g , the polynomial h has a single quadratic irreducible factor and no other irreducible factor of even degree.

- If $g \geq 2$, $\Omega_{4,\ell}$ is the set of $g \in CSp(2g, \mathbf{F}_\ell)$ with multiplier q such that if we factor $\det(1 - Tg)$ as before, the polynomial h has an irreducible factor of prime degree $> g/2$.

Indeed, spelling out again the relation between the factorization of P_t modulo ℓ and the existence of specific conjugacy classes in its Galois group, we have:

- (1) If P_t is *reducible*, then P_t can not be irreducible modulo ℓ (note that the leading term of P_t is $q^g T^{2g}$ and $\ell \neq p$ so the degree does not change by reduction), so by (8.19), this implies that $\rho_\ell(\text{Fr}_{t,q}) \notin \Omega_{1,\ell}$ for any $\ell \neq p$.
- (2) If P_t is irreducible but the Galois group G of its splitting field *does not contain a transposition* (when seen as a subgroup of \mathfrak{S}_{2g}), then $\rho_\ell(\text{Fr}_{t,q}) \notin \Omega_{2,\ell}$ for any ℓ : the opposite would imply that $P_t \pmod{\ell} = \det(1 - T\rho_\ell(\text{Fr}_{t,q}))$ has an irreducible factor of degree 2 and all others of odd degree, which means that G (still as a subgroup of \mathfrak{S}_{2g}) contains an element with cycle type consisting of one 2-cycle and further cycles of odd length, a power of which will be a transposition.

Next, if $g \geq 2$, notice that if we write (as in (8.22))

$$P_t = T^g Q_t(qT + T^{-1})$$

for a unique monic polynomial $Q_t \in \mathbf{F}_\ell[T]$ of degree g , the cycle in \mathfrak{S}_g associated to the polynomial Q_t is the image by the map $p : W_{2g} \rightarrow \mathfrak{S}_g$ of the cycle type associated to P_t . Indeed, because disjoint cycles are involved, it suffices to check that if P_t is irreducible, then so is Q_t , which is clear by contraposition.

We deduce from this and the same reasoning used in (2) that:

- (3) If P_t is irreducible but $p(G)$ *does not contain a transposition*, then we have $\rho_\ell(\text{Fr}_{t,q}) \notin \Omega_{3,\ell}$ for any ℓ .
- (4) If P_t is irreducible but $p(G)$ *does not contain a cycle of prime order* $m > g/2$, then $\rho_\ell(\text{Fr}_{t,q}) \notin \Omega_{4,\ell}$ for any ℓ .

On the other hand, if none of these four (or two if $g = 1$) possibilities hold, then Exercise 8.1 shows that $G = W_{2g}$. In other words, the exceptional set

¹² As we can in a unique way, because of the functional equation of characteristic polynomials of symplectic similitudes.

$N(f) \subset U(\mathbf{F}_q)$ satisfies

$$N(f) \subset \bigcup_{1 \leq i \leq 4} S(U(\mathbf{F}_q), \Omega_i; \mathcal{L}^*)$$

and hence

$$|N(f)| \leq \left(q + (2g-1)L^{2g^2+g+2}\sqrt{q} \right) (H_1^{-1} + H_2^{-1} + H_3^{-1} + H_4^{-1}) \quad (8.23)$$

by the large sieve, with obvious notation (and the last two terms should be omitted if $g = 1$).

Each of the terms H_i is of the type

$$\sum_{\substack{\psi(m) \leq L}}^{\flat} \beta_i(m)$$

where β_i is a multiplicative function, namely

$$\beta_i(m) = \prod_{\substack{\ell|m \\ (m, 2p)=1}} \frac{|\Omega_{i,\ell}|}{|Sp(2g, \mathbf{F}_\ell)| - |\Omega_{i,\ell}|}.$$

Moreover, the function $\beta_i(m)$ is well understood for $m = \ell$ a prime: by the results of Appendix B, we have

$$\beta_i(\ell) = \frac{|\Omega_{i,\ell}|}{|Sp(2g, \mathbf{F}_\ell)| - |\Omega_{i,\ell}|} = \frac{\delta_i}{1 - \delta_i} + o\left(\frac{1}{\ell}\right)$$

for $\ell \geq 3$, $\ell \neq p$, where δ_i is the density of the set of conjugacy classes in \mathfrak{S}_g or \mathfrak{S}_{2g} associated to the type of factorization permitted for the relevant polynomials, namely:

- $\delta_1 = (2g)^{-1}$ is the density of $2g$ -cycles in \mathfrak{S}_{2g} ;
- δ_2 is the density of elements in \mathfrak{S}_{2g} which are products of one transposition and disjoint cycles of odd length; one can easily check that $\delta_2 \geq (4g)^{-1}$;
- δ_3 is the same density in \mathfrak{S}_g as that of the set of conjugacy classes called C_1 in the proof of Theorem 4.2; by (4.3), we have $\delta_3 \sim (\log 2)/(\log g)$ as $g \rightarrow +\infty$, and clearly $\delta_3 > 0$ for all $g \geq 2$;
- similarly, δ_4 is the density in \mathfrak{S}_g of the set of conjugacy classes C_2 in the proof of Theorem 4.2; by (4.4), we have $\delta_4 \sim 1/\sqrt{2\pi g}$ as $g \rightarrow +\infty$, and also clearly $\delta_4 > 0$ for all $g \geq 2$.

From this, we see that we can apply Theorem G.2 from Appendix G, which is due to Lau and Wu, with $f(m) = \beta_i(m)$, $g(m) = \psi(m)$, the parameters being

$$(\kappa, \eta, \alpha, \alpha', \theta') = (\delta_i/(1 - \delta_i), 1, 1, 1, 0)$$

and in this manner we obtain in particular

$$H_i = \sum_{\psi(m) \leq L}^b \beta_i(m) \gg L(\log L)^{-1+\delta_i/(1-\delta_i)}, \quad (8.24)$$

for $L \geq 3$, where the implied constant depends only on g (we do not need the asymptotic formula here).

If we now select L such that $L^{2g^2+g+2} = q^{1/2}$ (if this leads to a value ≥ 3), the upper bound for $|N(f)|$ of the theorem follows. As usual, if this value of L is < 3 , we merely remark that by enlarging the implied constant, the statement is trivial. \square

Remark 8.16 In proving Theorem 8.15, we concentrated on a fixed genus g and tried to obtain the sharpest possible result. However, one can obtain weaker results more easily, e.g., using only the prime sieve support \mathcal{L}^* one gets

$$|\{t \in \mathbf{F}_q \mid f(t) \neq 0 \text{ and } [K_t : \mathbf{Q}] < 2^g g!\}| \ll q^{1-\gamma}(\log q)$$

for $L \geq 2$, the implied constant depending only on g , still with $\gamma = (4g^2 + 2g + 4)^{-1}$ but without needing the delicate results on sums of multiplicative functions of Appendix G. One can also obtain a result which is uniform in g , using the precise explicit bounds of Appendix B for the number of symplectic matrices with a given reversed characteristic polynomial. This was done in [80, Theorem 6.2]: one gets the same bound

$$|\{t \in \mathbf{F}_q \mid f(t) \neq 0 \text{ and } [K_t : \mathbf{Q}] < 2^g g!\}| \ll g^2 q^{1-\gamma}(\log q) \quad (8.25)$$

with an *absolute* implied constant (in [80] the gain γ is smaller, but this comes from using weaker estimates for the dimensions of irreducible representations of $CSp(2g, \mathbf{F}_\ell)$ than those of Chapter 5; also the factor g^2 is missing because of a small slip in handling the final estimates).

It is likely that one can refine Theorem 8.15 to make it also uniform in g , which amounts to checking the dependency on g in the estimates (8.24) for sums of multiplicative functions. However, the gain compared to (8.25) is of size $(\log q)^\delta$ with $\delta \sim 1/4g$, and thus becomes trivial as soon as g is of size $\log \log q$, while (8.25) is non-trivial for g somewhat smaller than $\sqrt{(\log q)}$ – already, a fairly short range.

Obtaining uniform estimates in terms of g should not be thought of as being simply an academic problem. Indeed, the applications of the sieve for Frobenius to families of zeta functions (or L -functions more generally) are also relevant to the delicate issues surrounding the use of Random Matrix Theory to investigate the arithmetic properties of L -functions *over number fields* (see the introduction to [77] for a survey of the problems involved). The ‘random matrices’ we have

here are precisely those of the Frobenius acting on $H^1(C_t \times \bar{\mathbf{F}}_q, \mathbf{Q}_\ell)$. In that context, the most important limit is that when the *size* of the matrices becomes large, which means taking $g \rightarrow +\infty$.

The maximality of the splitting field for a given curve has some interesting consequences. It may be interpreted as stating that the zeros of the zeta function are 'as independent as possible', and one can deduce various statements along those lines. In [84], it is shown that together with an ordinarity assumption (which can be phrased as asking that the coefficient of T^g in $P(T)$ is coprime with p), this maximality implies that the multiplicative subgroup of \mathbf{C}^\times generated by the roots of the zeta function is a free abelian group of rank $g + 1$. One can also show fairly easily that the maximality property and the additional condition that $\text{Tr}(\text{Fr} \mid H_c^1(C \times \bar{\mathbf{F}}_q, \mathbf{Q}_\ell)) \neq 0$ imply that the zeros of the zeta function are \mathbf{Q} -linearly independent (see Exercise 8.2 below). Obviously, a further use of sieve or individual equidistribution will show that this additional condition holds for most parameters in the families of hyperelliptic curves considered previously.

Such results are of interest as analogues of conjectures concerning the linear or algebraic independence of roots of L -functions over number fields. These conjectures have appeared in a number of investigations (see, for instance the treatment by Rubinstein and Sarnak [109] of the 'Chebychev bias' among the residue classes of primes to a given modulus, where the hypothesis that the non-negative ordinates of zeros of primitive Dirichlet L -functions are \mathbf{Q} -linearly independent plays an important role).

Remark 8.17 The method should not be thought of as depending intrinsically on the 'large monodromy' assumption. For instance, if one happened to know that the geometric monodromy group is, for most ℓ , of the type $Sp(2g_1, \mathbf{F}_\ell) \times Sp(2g_2, \mathbf{F}_\ell)$ with g_1, g_2 fixed positive integers, one would expect to show that it follows that for most parameters, the numerator of the zeta function factors over \mathbf{Q} as a product of two irreducible polynomials, of degree $2g_1$ and $2g_2$ respectively. More interesting are cases with orthogonal monodromy (i.e., G_ℓ^θ an orthogonal group or a special orthogonal group for a non-degenerate symmetric bilinear form), where there are sometimes forced eigenvalues, depending on the parity of the dimension of the relevant space and the determinant of the orthogonal matrix. This is arithmetically very relevant in terms of 'trivial' central zeros of L -functions. Then one would wish to show that, after dividing by the obvious factor of the characteristic polynomial, what remains is irreducible. Katz [75] has proved results of this type similar to Chavdarov's original work (see also [76] for earlier result with de Jong); F. Jouve is currently adapting the large sieve to this situation.

Exercise 8.2 This exercise shows that if C/\mathbf{F}_q is a smooth projective geometrically irreducible algebraic curve over a finite field such that the splitting field of the numerator $P(T)$ of the zeta function of C is W_{2g} , then the roots of $P(T)$ are \mathbf{Q} -linearly independent unless

$$\mathrm{Tr}(\mathrm{Fr} \mid H_c^1(C \times \overline{\mathbf{F}}_q, \mathbf{Q}_\ell)) = 0. \quad (8.26)$$

For this, we use methods of Girstmair [48] that can be used more generally to classify the linear (or polynomial) relations between roots of a polynomial over a field. Let $\beta_1, \dots, \beta_{2g}$ be the roots of P in \mathbf{C} , in pairs $(\beta_{2i-1}, \beta_{2i})$ with $\beta_{2i-1}\beta_{2i} = q$.

- (1) Show that W_{2g} acts \mathbf{Q} -linearly on the \mathbf{Q} -vector space E generated by the β_i , and on the free vector space F generated by symbols $[\beta_i]$ for each root.
- (2) Show that the \mathbf{Q} -vector space

$$R = \left\{ (\lambda_1, \dots, \lambda_{2g}) \in \mathbf{Q}^{2g} \mid \sum_{i=1}^{2g} \lambda_i \beta_i = 0 \right\}$$

may be identified with the kernel of the obvious W_{2g} -linear map $F \rightarrow E$.

- (3) Show that F , as a representation of W_{2g} , is isomorphic to $\mathrm{Ind}_H^G(1)$, where H is the stabilizer of β_1 in W_{2g} (seen as acting on the roots). [**Hint:** Use the bijection between W_{2g}/H and the roots β_i of P to see F as a permutation representation of W_{2g}/H .]
- (4) Deduce that F decomposes as the direct sum of irreducible representations

$$F = F_0 \oplus F_1 \oplus F_2$$

where F_0 is the trivial component, of dimension one generated by $\sum [\beta_i]$, and

$$F_1 = \left\{ \sum_{i=1}^{2g} \lambda_i [\beta_i] \mid \lambda_{2i-1} - \lambda_{2i} = 0, \ 1 \leq i \leq g, \sum \lambda_i = 0 \right\},$$

$$F_2 = \left\{ \sum_{i=1}^{2g} \lambda_i [\beta_i] \mid \lambda_{2i-1} + \lambda_{2i} = 0, \ 1 \leq i \leq g \right\}.$$

[**Hint:** Show that $F \otimes \mathbf{C}$ is the sum of three irreducible representations of W_{2g} (one can use for instance [115, Exercise 2.6] and check that there are three orbits of W_{2g} acting on $W_{2g}/H \times W_{2g}/H$).]

- (5) By making a list of possibilities for the subrepresentation R , show that only $R = 0$ or $R = F_0$ are possible, and the latter is equivalent with (8.26).

[Hint: For instance, if R contains F_1 , then R contains $\sigma(\beta_1) - \beta_1 + \sigma(\beta_2) - \beta_2$ for all $\sigma \in W_{2g}$, giving $\beta_1 + \beta_2 \in \mathbf{Q} \dots$]

Of course, it would be quite interesting to have examples of *lower bounds* for the number of parameters t where the polynomial P_t does not have a maximal splitting field. In particular, it is not clear at all if (under the conditions of this theorem) the set of $t \in \overline{\mathbf{F}}_q$ where this holds is infinite – we have no example one way or another. Still, as in Section 7.7, some numerical experiments are possible. Note however that this possibility is a very recent development, depending on the discovery and implementation of efficient algorithms for computations of zeta functions of hyperelliptic curves over finite fields. Specifically, we used a recent algorithm of Hubrechts [62] in MAGMA 2.13, which is based on p -adic techniques and a mixture of other recent ideas of Kedlaya and Lauder (this technique is especially well-suited for our purposes since it is adapted to computations of zeta functions for families of curves, dealing simultaneously with many values of t much faster than individually). The computations lasted a few days on a fast Opteron machine.

We first looked at the two families of curves of genus 3 given by

$$C_t : y^2 = (x^6 + x - 1)(x - t), \quad D_t : y^2 = (x^6 + x^3 - x - 1)(x - t)$$

over \mathbf{F}_5 (which were chosen ‘randomly’ by pure thought), and for those we computed all zeta functions over \mathbf{F}_{5^k} for degrees $k \leq 8$. For each degree, we computed which numerators are reducible, and furthermore which irreducible numerators have Galois group of order $< 48 = |W_6|$.

Since Galois-conjugate parameters (over \mathbf{F}_5) yield isomorphic curves, we give results listing only the number of ‘exceptional’ parameters t up to Galois conjugation. We also list the factorization type or the non-maximal Galois group. Precisely, the columns of the tables below are as follows:¹³

- the degree k of the parameters in the current row;
- the number of parameters of degree k with the factorization type or Galois group in the third column (up to conjugation);
- the factorization-type, where P_i denotes a polynomial of degree i , of the numerator of the zeta function of C_t , or the Galois group if it is irreducible with non-maximal splitting field.

Of course, one notices immediately in the case of C_t that many more examples occur in the field \mathbf{F}_{5^5} . This may be because the characteristic divides the degree,

¹³ The algorithm currently implemented in MAGMA is not applicable when $t = 0$; so the data omits this point and the results for degree 1 may be off by one.

Table 8.1 *Non-generic zeta functions for*
 $y^2 = (x^6 + x - 1)(x - t)$

Degree	Number	Factorization/Galois group
2	1	$P_2 P_4$
4	2	$P_2 P_4$
5	10	$P_2 P_4$
8	3	D_{12}

Table 8.2 *Non-generic zeta functions for*
 $y^2 = (x^6 + x^3 - x - 1)(x - t)$

Degree	Number	Factorization/Galois group
1	1	$P_2 P_4$
2	4	$P_2 P_4$
5	2	$P_2 P_4$
6	3	$P_2 P_4$
7	1	$P_2 P_4$
8	23	$P_2 P_4$
8	1	P_2^3

or because the polynomial $x^6 + x - 1$ factors as $(x - 2)(x^5 + 2x^4 - x^3 - 2x^2 + x - 2)$ in $\mathbf{F}_5[x]$. Also note that there are no parameters of degree 6, 7 or 8. For D_t , we have $x^6 + x^3 - x - 1 = (x + 1)(x + 4)(x^4 + x^2 + x + 1)$ in $\mathbf{F}_5[x]$, but there is no particular ‘spike’ of reducible parameters over \mathbf{F}_{5^4} . No examples of irreducible polynomials with non-maximal Galois groups were found in the second family.

Finally, we performed some computations using a family defined over the base field \mathbf{F}_{5^6} , defined by

$$C_t : y^2 = (x^6 - \omega x - 1)(x - t)$$

where ω is a generator of \mathbf{F}_{5^6} defined by the minimal polynomial $x^6 + x^4 - x^3 + x^2 + 2$. We computed the zeta functions for $t \in \mathbf{F}_{5^6}$, and found 3 values of t , of degree 6, for which the numerator of the zeta function factors as $P_2 P_4$, but no instances of small Galois group.

All in all, these experiments amount to computing roughly 160 000 zeta functions (counting parameters up to Galois-conjugacy; of course non-conjugate parameters may sometimes lead to the same curve), with only 51 cases of reducible polynomials and 3 occurrences of non-maximal Galois groups. More extensive experiments would certainly be quite useful, in particular involving higher-genus curves.

8.7 Remarks on monodromy groups

The proof of Theorem 8.15 hinges crucially on the computation of G_ℓ^g , which is given by Yu's theorem (Proposition 8.13). Indeed, it seems likely that most interesting applications of the sieve for Frobenius will depend on knowing quite precisely the geometric monodromy groups of the family (ρ_ℓ) .

This is a delicate issue in general, but here a few simple remarks. Assume that (ρ_ℓ) is obtained by reduction modulo ℓ from a family of representations $\tilde{\rho}_\ell : \pi_1(U, \bar{\eta}) \rightarrow GL(r, \mathbf{Z}_\ell)$. One can then also consider the images of $\pi_1(U, \bar{\eta})$ and $\pi_1(\bar{U}, \bar{\eta})$ by $\tilde{\rho}_\ell$, which are (compact) subgroups of $GL(r, \mathbf{Z}_\ell)$. Knowing these integral monodromy groups would be even better than knowing G_ℓ^g , but this is also harder. However, it is often *easier* to compute the Zariski closure $\tilde{G}_\ell^g \subset GL(r, \bar{\mathbf{Q}}_\ell)$ of $\tilde{\rho}_\ell(\pi_1(\bar{U}, \bar{\eta}))$. Recall that, by definition, this group is (or at least can be identified with) the smallest group of matrices $g \in GL(r, \bar{\mathbf{Q}}_\ell)$ such that, for any polynomial $P \in \bar{\mathbf{Q}}_\ell[X_{i,j}, D]$, with $1 \leq i, j \leq r$, we have

$$\tilde{P}(g) = 0$$

if $\tilde{P}(h) = 0$ for all $h \in \tilde{\rho}_\ell(\pi_1(\bar{U}, \bar{\eta}))$, where

$$\tilde{P}(g) = P(g_{i,j}, 1/\det(g))$$

for any matrix $g \in GL(r, \bar{\mathbf{Q}}_\ell)$. (In other words, this is the largest group of matrices which *can not be distinguished* from $\tilde{\rho}_\ell(\pi_1(\bar{U}, \bar{\eta}))$ using only polynomial functions of the coordinates and of the determinant.)

Why this group, which is called the *rational geometric monodromy group* of $\tilde{\rho}_\ell$, should be any easier to apprehend may seem a mystery at first; one point which is easy to make is that it is a 'continuous' object, in a sense, not a discrete one, and that continuous phenomena are often rather simpler than purely discrete ones. In fact, this group was shown to be of a rather special kind; for instance,¹⁴ provided the representation

$$\pi_1(\bar{U}, \bar{\eta}) \rightarrow GL(r, \bar{\mathbf{Q}}_\ell)$$

is *semisimple* (i.e., a direct sum of irreducible subrepresentations, which is not automatic here but holds, in particular, if it is irreducible), the geometric monodromy group has a faithful completely reducible linear representation, which implies that its connected component of the identity is a *reductive group*. The point is that such groups are quite rigid;¹⁵ see Appendix E for a quick survey

¹⁴ This is weaker than the known results.

¹⁵ Their classification, in particular, is essentially independent of ℓ , which is rather crucial to the philosophy according to which the monodromy group of a compatible system should also be independent of ℓ .

of the definitions of reductive linear algebraic groups. To give but an inkling of what this entails, this proves that it is *not possible* (under the conditions stated) that

$$\tilde{G}_\ell^g = \left\{ g = \begin{pmatrix} a & b \\ 0 & d \end{pmatrix} \mid a, b, d \in \overline{\mathbf{Q}}_\ell, ad \neq 0 \right\},$$

simply because such a group of matrices is not reductive. (On the other hand, it is perfectly possible for the *finite monodromy* to satisfy

$$G_\ell^g = \left\{ g = \begin{pmatrix} a & b \\ 0 & d \end{pmatrix} \mid a, b, d \in \mathbf{F}_\ell, ad \neq 0 \right\},$$

in particular cases.)

In the case of Yu's theorem, the analogue of Proposition 8.13 for \tilde{G}_ℓ^g is proved (essentially from scratch) by Katz–Sarnak in [77, Theorem 10.1.16]: we have

$$\tilde{G}_\ell^g = Sp(2g, \overline{\mathbf{Q}}_\ell) \quad (8.27)$$

for all primes ℓ (including $\ell = 2$).

Now we could derive Theorem 8.15 from this fact, instead of appealing to Yu's theorem, using a remarkable result of Larsen [85, Theorem 3.17], which shows that (8.27) for all ℓ , *for a compatible system* of representations $(\tilde{\rho}_\ell)$, implies that

$$G_\ell^g = Sp(2g, \mathbf{F}_\ell)$$

for a set of primes of density 1, i.e., for all $\ell \in \Lambda$, where Λ is such that

$$\lim_{L \rightarrow +\infty} \frac{|\{\ell \in \Lambda \mid \ell \leq L\}|}{\pi(L)} = 1.$$

While this may look like a simpler approach to the solution of Chavdarov's problem, there are two issues to keep in mind (in this particular situation, where an alternative exists):

- Larsen's theorem involves the classification of finite simple groups; although it uses it in a robust way (i.e., finding finitely many new exceptional finite simple groups would not affect the argument at all, and any infinite family that 'behaves' like those already known could certainly be handled without trouble), this still introduces a dependency¹⁶ on such a vast body of knowledge that it is hard to resist feeling that one's work becomes the mere addition of a footnote to the theory of finite groups.
- The set of primes Λ given by Larsen's theorem is not explicit (so, even though we know it is very large, there is no way to say what is the smallest prime ℓ to

¹⁶ The work of Yu doesn't use the classification, and neither does the alternative proof by C. Hall, although it appeals to non-trivial results concerning finite groups (due to Zalesskiĭ and Serežkin).

which it applies); this means that it is not possible to use it to prove uniform results when it is applied infinitely many times; e.g., it can not be used to prove (8.25) for all g , with an absolute implied constant.

There remain cases, however, where Larsen's result is the only way to prove the desired result, and in particular it provides a quick solution to sieve problems whenever the rational monodromy groups are known. This, it turns out, is quite often the case, due especially to the many pioneering works of Katz (see [72], [73], for instance). Moreover, the latest work of Katz [74], involving the so-called 'Larsen alternative', provides new criteria, of a very arithmetic nature, to (almost) determine the rational monodromy group based on what seems like magically little information!

A last useful remark is that in order to show that the (rational or finite) geometric monodromy group of a family (ρ_t) on a parameter variety U/\mathbf{F}_q is as large as possible, it suffices to show that this is so for a subvariety (intuitively, we are just saying that if a subfamily 'varies maximally' then so does the full family, which is quite natural). Even more generally, we state this in an obvious lemma:

Lemma 8.18 *Let U/\mathbf{F}_q be a smooth geometrically connected affine variety, and $\pi_1(\overline{U}, \bar{\eta}) \rightarrow G$ a continuous homomorphism to a finite group G . Let*

$$\overline{V} \xrightarrow{f} \overline{U}$$

be an arbitrary morphism from another smooth geometrically connected variety \overline{V} over \mathbf{F}_q . If, for some geometric generic point $\bar{\eta}'$ of \overline{V} mapping to $\bar{\eta}$, the representation

$$\pi_1(\overline{V}, \bar{\eta}') \rightarrow \pi_1(\overline{U}, \bar{\eta}) \rightarrow G$$

is onto, then the same holds for the original homomorphism.

Here is an application, where the 'subfamily' is given by one of the one-parameter families of Example 8.6.

Proposition 8.19 *Let $q = p^k$ and $g \geq 2$ such that $p > 2g + 1$. Then the number $N(g, q)$ of isomorphism classes¹⁷ of smooth projective geometrically irreducible curves C/\mathbf{F}_q such that the numerator of the zeta function of C is*

¹⁷ Isomorphism of (smooth projective geometrically connected) algebraic curves C/\mathbf{F}_q can be seen as isomorphism of their function fields $\mathbf{F}_q(C)$, which are algebraic extensions of finite degree of the field $\mathbf{F}(T)$ of rational functions over \mathbf{F}_q .

either reducible or has splitting field with Galois group strictly smaller than W_{2g} satisfies

$$N(g, q) \ll q^{3g-3-\gamma} (\log q)$$

where $\gamma = 1/(12g^2 + 6g + 8)$, and the implied constant depends only on p and g .

Proof The idea is to use an algebraic parameter space U/\mathbf{F}_q (called classically a ‘moduli space’) which classifies the isomorphism classes of curves C/\mathbf{F}_q , i.e., such that each $u \in U(\mathbf{F}_q)$ corresponds to a unique curve C/\mathbf{F}_q . Although it is well known that this is not possible in a strict sense (because of problems with curves having automorphisms), algebraic geometers have found various ways to work around this difficulty. The simplest technique is to use moduli spaces which classify curves (C, r) with additional ‘rigidifying’ data r , so that any enriched curve (C, r) has trivial automorphism group. There is a precise and enlightening discussion of this in Chapter 10 of [77]; following Sections 10.5 and 10.6 of [77], we use a moduli space U_g with the following property: U_g/\mathbf{F}_q is a smooth affine geometrically connected algebraic variety of dimension

$$\dim U_g = 3g - 3 + (5g - 5)^2$$

such that for any $n \geq 1$, there is a natural bijection

$$U(\mathbf{F}_{q^n}) \simeq \{(C, r)\} / \sim \quad (8.28)$$

where the pairs (C, r) consist of a smooth projective geometrically connected algebraic curve U/\mathbf{F}_{q^n} of genus g , together with a basis $r = (r_1, \dots, r_{5g-5})$ of the \mathbf{F}_{q^n} -vector space $\Gamma(C, \Omega^1)^{\otimes 3}$ (the third tensor power of the vector space of 1-differentials on C which are everywhere defined; this is of dimension $5g - 5$ by the Riemann–Roch formula), and the equivalence relation on pairs is the ‘obvious’ notion of isomorphism: $(C_1, r_1) \sim (C_2, r_2)$ if and only if there is an isomorphism $f : C_1 \rightarrow C_2$ (as algebraic curves) such that $f(r_1) = r_2$.

Note that these properties, and in particular the irreducibility of U_g , are highly non-trivial facts of algebraic geometry, due to Deligne and Mumford among others (see again the discussion in Section 10.6 of [77] for detailed references; what we denote by U_g is denoted by $\mathcal{M}_{g,3K}$ there).

On U_g there is a ‘tautological’ algebraic family of curves (with additional structure)

$$\mathcal{C}_g \xrightarrow{\pi} U_g$$

such that the fiber over a point $u \in U(\mathbf{F}_{q^n})$ is precisely $C \times \{r\}$, the curve C/\mathbf{F}_{q^n} associated to u by the bijection (8.28) with all the bases of the vector

space $\Gamma(C, \Omega^1)^{\otimes 3}$. We apply the sieve for Frobenius to U_g with the family of reductions modulo ℓ of the compatible system $R^1\pi_*\mathbf{Z}_\ell$.

Now, select (arbitrarily) a monic polynomial $f \in \mathbf{F}_{q^n}[T]$ (for some $n \geq 1$)¹⁸ which is squarefree and of degree $2g$. The existence of the algebraic family of curves

$$\mathcal{C} \xrightarrow{\pi'} V_f$$

with equations

$$y^2 = f(x)(x - t)$$

(i.e., those obtained by compactification and desingularization, as in Example 8.6), parametrized by the open subspace V_f of the affine line over \mathbf{F}_{q^n} where f does not vanish, can be shown (see the proof of Theorem 10.6.11 in [77] for the fact that the family of curves above over V_f can be lifted to an algebraic family of curves with the additional $3K$ -structure) to imply that there exists a morphism $V_f \rightarrow U_g$ such that the composition

$$\pi_1(\overline{V}_f, \bar{\eta}') \rightarrow \pi_1(\overline{U}_g, \bar{\eta}) \rightarrow Sp(2g, \mathbf{F}_\ell)$$

(for some suitable $\bar{\eta}'$) ‘is’ the representation of $\pi_1(\overline{V}_f, \bar{\eta}')$ associated with $R^1\pi_*\mathbf{F}_\ell$. So by Yu’s theorem, Lemma 8.18 applies with $G = Sp(2g, \mathbf{F}_\ell)$ for each $\ell \nmid 2p$, showing that the finite geometric monodromy group for ρ_ℓ is $Sp(2g, \mathbf{F}_\ell)$ for all $\ell \nmid 2p$. (In fact, this holds for $p = 2$ also, though this does not follow from Yu’s theorem; see [4].)

Since the dimension of the parameter space is > 1 , we must use Case (i) of the sieve for Frobenius (Corollary 8.10), and in particular ensure that the family (ρ_ℓ) is restricted to a set of primes Λ for which the action is tame. For this purpose, notice that if $r \geq 1$ is an integer and $p > r + 1$ a prime number, there exists $\alpha \in (\mathbf{Z}/p\mathbf{Z})^\times$ such that the order of $GL(r, \mathbf{F}_\ell)$ is prime to p if ℓ satisfies $\ell \equiv \alpha \pmod{p}$. (Indeed, from (0.1), we see that the order

$$\ell^{r(r-1)/2} \prod_{1 \leq i \leq r} (\ell^i - 1)$$

of $GL(r, \mathbf{F}_\ell)$ is prime to p if $p \equiv \alpha \pmod{\ell}$ whenever the order of α modulo p is $> r$; if $p > r + 1$, any primitive root α modulo p will certainly do.) Since $p > 2g + 1$ by assumption, we can select such an α for $r = 2g$, and consider the set Λ of odd prime numbers $\ell \equiv \alpha \pmod{p}$; then the geometric monodromy group is of order prime to p for $\ell \in \Lambda$ as a subgroup of $GL(2g, \mathbf{F}_\ell)$.

Now we apply the sieve for Frobenius with the same sieving sets as in Theorem 8.15; for simplicity, consider only $\mathcal{L} = \mathcal{L}^*$ where \mathcal{L}^* is the set of primes

¹⁸ Note that we can take $n = 1$ here because we assume $p > 2g + 1$, but if we try to extend the proposition to all p and g , this may require taking $n \neq 1$.

in $\Lambda \leq L$ (the assumptions of Theorem G.2 are not satisfied for a sum over integers divisible only by a sparse sequence of primes). Everything goes through with the lower bound

$$\sum_{\substack{3 < \ell \leq L \\ \ell \in \Lambda}} \frac{|\Omega_{i,\ell}|}{|Sp(2g, \mathbf{F}_\ell)|} \geq \delta_i \pi(L) + O(1)$$

for $i = 1, \dots, 4$ and $L \geq 2$, the implied constant depending on p and g . This provides the bound

$$\tilde{N}(g, q) \ll q^{\dim U_g - \gamma} (\log q)$$

where¹⁹ $\gamma = 12g^2 + 6g + 8$, the implied constant depending on p and g , where $\tilde{N}(g, q)$ is the number of pairs $(C, r) \in U(\mathbf{F}_q)$ where the splitting field of the numerator of the zeta function of C is small.

Now, notice that for any pair (C, r) which is counted in $\tilde{N}(g, q)$, and for any $x \in GL(\Gamma(C, \Omega^1)^{\otimes 3}) \simeq GL(5g - 5, \mathbf{F}_q)$, the pair $(C, x \cdot r)$ is also counted. Moreover, there are at most $|\text{Aut}(C)|$ such pairs $(C, x \cdot r)$ which give the same point $u \in U(\mathbf{F}_q)$, by definition of the equivalence relation. The size of the automorphism group is bounded in terms of g only (see, e.g., [77, Lemma 10.6.12]), say $|\text{Aut}(C)| \leq \beta_g$, and it follows that

$$\begin{aligned} N(g, q) &\leq \beta_g \tilde{N}(g, q) |GL(5g - 5, \mathbf{F}_q)|^{-1} \\ &\ll q^{3g-3+(5g-5)^2-\gamma} q^{-(5g-5)^2} (\log q) = q^{3g-3-\gamma} (\log q) \end{aligned}$$

where the implied constant depends only on p and g , as desired. \square

Problem 8.20 It remains an open question to extend this proposition to curves of all genus $g \geq 1$ over finite fields of all characteristics.

Remark 8.21 A recent paper of Achter and Pries [4] shows that the geometric monodromy group of the p -rank strata of the moduli space of curves of genus $g \geq 1$ is still $Sp(2g, \mathbf{F}_\ell)$ for all $\ell \neq p$, with the exception of the supersingular stratum of curves of genus ≤ 2 .²⁰ So, as stated in [4], Proposition 8.19 extends to curves with a specified p -rank $f \in \{0, \dots, g\}$, with $f \geq 1$ if $g \leq 2$.

¹⁹ This constant is $2A$ where $A = v + 2s + t + 1$ for $G_\ell \subset CSp(2g, \mathbf{F}_\ell)$, see Section 8.5.

²⁰ The p -rank is related to p -adic properties of the eigenvalues of Frobenius; for instance, maximal p -rank corresponds to *ordinary* curves; those where, among all pairs $(\alpha, q/\alpha)$ of eigenvalues, one of the two is coprime with p (in the ring of all algebraic integers). This stratum is dense in the moduli space.

8.8 A last application

We conclude this chapter, and the main part of the book, with a proof of Theorem 1.6, which we recall from the introduction:

Theorem 8.22 *Let q be a power of a prime number $p \geq 5$, $g \geq 1$ an integer and let $f \in \mathbf{F}_q[T]$ be a squarefree polynomial of degree $2g$. For t not a zero of f , let C_t denote the smooth projective model of the hyperelliptic curve*

$$y^2 = f(x)(x - t),$$

and let J_t denote its Jacobian variety.²¹ Then we have

$$|\{t \in \mathbf{F}_q \mid f(t) \neq 0 \text{ and } |C_t(\mathbf{F}_q)| \text{ is a square}\}| \ll gq^{1-\gamma}(\log q),$$

$$|\{t \in \mathbf{F}_q \mid f(t) \neq 0 \text{ and } |J_t(\mathbf{F}_q)| \text{ is a square}\}| \ll gq^{1-\gamma}(\log q)$$

where $\gamma = (4g^2 + 2g + 4)^{-1}$, and the implied constants are absolute.

This result is only a small (interesting) step along the way for the general problem (still badly understood) of the arithmetic properties of the number of points of algebraic varieties over finite fields, compared with ‘random’ integers; see also the end of Appendix A for a lower bound sieve result on the same families of curves. These questions have become quite important because their answers have direct consequences concerning the performance of some important algorithms based on properties of algebraic varieties over finite fields, e.g., elliptic curve factorization and primality testing (introduced by H. Lenstra), and elliptic curve public key cryptography (introduced by Koblitz and Miller), as well as their generalizations to curves of higher genus (in particular hyperelliptic curves, of which the families of curves above are examples; elliptic curves correspond to $g = 1$). Note that for factorization and primality testing, one needs the number of points to be *friable* integers, i.e., divisible by (many) small primes,²² whereas for cryptographic applications, one wants the number of points to be essentially prime.

The only real issue in this result, for analytic number theorists at least, might be to recall what is the Jacobian $J(C)$ of a curve C . We do this in a few words, for the special case of a curve over a finite field. See, e.g., [90, 7.4.4] for more detailed information; the actual *existence* of the Jacobian is again a deep result of algebraic geometry (due to Weil in the case of a curve over an arbitrary field). Let C/k be a smooth projective geometrically connected algebraic curve over

²¹ See below for a few words of explanation if this is not a familiar notion.

²² Those numbers are rather misleadingly called ‘smooth’ in much of the non-French literature.

a finite field k . Then there exists a smooth projective variety $J(C)$, defined over the same field k , of dimension equal to the genus g of C , which has the following property: the set $J(C)(\bar{k})$ of points of J with coordinates in an algebraic closure of k is naturally in bijection with the abelian group $\text{Pic}^0(\bar{C}) = \text{Div}^0(\bar{C})/P(\bar{C})$ of classes of divisors of degree 0 on \bar{C} (i.e., the curve C seen ‘geometrically’ over the algebraic closure of k) modulo the subgroup of principal divisors. In other words, this is the quotient of the group of formal integral linear combinations of points in $C(\bar{k})$ of the form

$$D = \sum_{x \in C(\bar{k})} a(x)[x],$$

such that the degree

$$\deg(D) = \sum_{x \in C(\bar{k})} a(x)$$

is zero, modulo the subgroup of *principal* divisors, of the type

$$(f) = \sum_{x \in C(\bar{k})} \deg(x)$$

for a non-zero rational function $f \in \bar{k}(C)$, where $\text{ord}_x(f)$ is the order of the zero (or pole if negative) of the rational function f at x ; the fact that $\deg(f) = 0$ reflects the property that a non-zero rational function has as many zeros as poles, counted with multiplicity.

We can define the zeta function of $J(C)$ by the same formula as (8.2):

$$Z(J(C), T) = \exp \left(\sum_{n \geq 1} |J(C)(\mathbf{F}_{q^n})| \frac{T^n}{n} \right)$$

(where the points with coordinates in \mathbf{F}_{q^n} can be recovered from the above as those in $J(C)(\bar{k})$ invariant under the Galois group of $\mathbf{F}_{q^n}/\mathbf{F}_q$, i.e., under the n -th power of the Frobenius automorphism, which acts in the obvious way on divisors through its action on $C(\bar{k})$), and one then shows that it has an expression as a rational function

$$Z(J(C), T) = \frac{P_1(T)P_3(T) \cdots P_{2g-1}(T)}{P_0(T)P_2(T) \cdots P_{2g}(T)}$$

where $P_i(T)$ is a polynomial with integer coefficients, such that $P_0 = 1 - T$ and $P_{2g} = 1 - q^g T$, in particular. In fact, a cohomological expression similar to (8.18) exists for all i , $0 \leq i \leq 2g$, and states that

$$P_i(T) = \det(1 - T \text{Fr} \mid \bigwedge^i H_c^1(C \times \bar{\mathbf{F}}_q, \mathbf{Z}_\ell))$$

for any prime $\ell \neq p$, i.e., P_i is the reversed characteristic polynomial of the geometric Frobenius automorphism acting on the i -th exterior power of the first cohomology groups of $C \times \overline{\mathbf{F}}_q$. In particular, $P_1(T)$ is the same as the numerator of the zeta function of C itself (see (8.18) again).

Proof We can certainly afford to be rather brief here, since all ingredients have already been mentioned with a fair amount of detail (in fact, the proof could be considered as an exercise for many readers).

The sieve setting and siftable set are the same as in Theorem 8.15; the point is of course that the family (ρ_ℓ) already used provides a way to understand the number of points of C_i and J_i over \mathbf{F}_q . Indeed, they are given by the formulas

$$\begin{aligned} |C_i(\mathbf{F}_q)| &= q + 1 - \text{Tr}(\text{Fr} \mid H^1(\bar{C}_i, \mathbf{Z}_\ell)), \\ |J_i(\mathbf{F}_q)| &= |\det(1 - \text{Fr} \mid H^1(\bar{C}_i, \mathbf{Z}_\ell))|, \end{aligned}$$

for any prime $\ell \nmid p$. Both follow from the generating series definition of the zeta functions by comparing with their cohomological expressions (see (8.2) and (8.18) for the first one, and for the second remember that

$$\det(1 - XT \mid M) = \sum_{i=0}^r (-1)^i \text{Tr}(T \mid \bigwedge^i M) X^i$$

for any endomorphism T of a free module M of finite rank r over a ring).

Thus, defining sieving sets

$$\begin{aligned} \Omega_\ell^J &= \{g \in \text{CSp}(2g, \mathbf{F}_\ell) \mid m(g) = q, \text{ and } \det(g - 1) \text{ is a square in } \mathbf{F}_\ell\}, \\ \Omega_\ell^C &= \{g \in \text{CSp}(2g, \mathbf{F}_\ell) \mid m(g) = q, \text{ and } q + 1 - \text{Tr}(g) \text{ is a square in } \mathbf{F}_\ell\} \end{aligned}$$

(recall that $m(g)$ is the multiplier of a symplectic similitude), we have inclusions

$$\{t \in \mathbf{F}_q \mid f(t) \neq 0 \text{ and } |\mathbf{S}_t(\mathbf{F}_q)| \text{ is a square}\} \subset S(U(\mathbf{F}_q), \Omega^S; \mathcal{L}^*),$$

for $\mathbf{S} \in \{C, J\}$, valid for any prime sieve support \mathcal{L}^* .

By (3) and (4), respectively, of Proposition B.4 in Appendix B, we have

$$\frac{|\Omega_\ell^S|}{|\text{Sp}(2g, \mathbf{F}_\ell)|} \geq \frac{1}{2} \left(\frac{\ell}{\ell + 1} \right)^{2g^2 + g + 1},$$

for $\ell \geq 3$. Thus if \mathcal{L} is the set of odd primes $\leq L$, we obtain

$$|\{t \in \mathbf{F}_q \mid f(t) \neq 0 \text{ and } |\mathbf{S}_t(\mathbf{F}_q)| \text{ is a square}\}| \leq (q + (2g - 1)\sqrt{q}L^A)H^{-1}$$

where $A = 2g^2 + g + 2$, and

$$H = \sum_{3 \leq \ell \leq L} \frac{|\Omega_\ell^S|}{|Sp(2g, \mathbf{F}_\ell)|} \geq \frac{1}{2} \sum_{3 \leq \ell \leq L} \left(\frac{\ell}{\ell+1} \right)^{2g^2+g+1}.$$

By the mean-value theorem we have

$$\left(\frac{\ell}{\ell+1} \right)^{2g^2+g+1} = 1 + O\left(\frac{g^2}{\ell+1} \right)$$

for $\ell \geq 3$, $g \geq 1$, with an absolute implied constant, and thus by the Prime Number Theorem we have

$$H \geq \frac{1}{2} \pi(L) + O(g^2 \log \log L)$$

with an absolute implied constant. For $L \gg g^2 \log 2g \log \log 3g$, this gives

$$H \gg \frac{L}{\log L}$$

with an absolute implied constant, and therefore

$$|\{t \in \mathbf{F}_q \mid f(t) \neq 0 \text{ and } |S_t(\mathbf{F}_q)| \text{ is a square}\}| \ll g(q + q^{1/2} L^A) L^{-1} (\log 2L),$$

with an absolute implied constant. In fact, this last inequality holds for all $g \geq 1$ and $L \geq 1$, being trivial (for a sufficiently large implied constant) if $L \ll g \log 2g$, and a fortiori if $L \ll g^2 \log 2g \log \log 3g$. (Note that it would not hold with $\log 2L$ replaced by $\log L$, as L close to 1 would create a problem, and indeed when g is large compared with q , L will be very close to 1.)

Now we select $L = q^{1/(2A)}$ as usual, and we obtain the uniform estimate

$$|\{t \in \mathbf{F}_q \mid f(t) \neq 0 \text{ and } |S_t(\mathbf{F}_q)| \text{ is a square}\}| \ll gq^{1-\gamma} (\log q)$$

with $\gamma = 1/(4g^2 + 2g + 4)$, and with an absolute implied constant. □

Appendix A

Small sieves

A.1 General results

If we are in a general sieving situation as described in Chapter 2, we may in many cases be interested in a lower bound for the size (measure) of $S(X, \Omega; \mathcal{L}^*)$, in addition to the upper bounds that the large sieve naturally provides. For this purpose, we can hope to appeal to the usual principles of small sieves, at least when Λ is the set of prime numbers and for some specific sieve supports. We describe this for completeness, with no claim to originality, and refer to books such as [55], the forthcoming ‘Sieve Theory’ by H. Iwaniec and J. Friedlander, or [67, Section 6] for more detailed coverage of the principles of sieve theory. The results of Gamburd, Bourgain and Sarnak [14, 15] concerning orbits of discrete group actions are recent examples of applications of small sieves in a sophisticated context.

We assume that our sieve setting is of the type $\Psi = (Y, \{\text{primes}\}, (\rho_\ell))$, and our prime sieve support will be a set \mathcal{L}^* of prime numbers $\ell < L$ for some parameter L . The siftable set is of the general type (X, μ, F) , as in Chapter 2, and we write $S(X, \Omega; L)$ for the sifted set $S(X, \Omega; \mathcal{L}^*)$.

The two small sieve techniques which are most commonly used are the Selberg (or Λ^2) sieve and the combinatorial sieves of Brun–Iwaniec–Rosser type. We present the latter here, with just a few words concerning the former.

Let

$$\Omega_d = \prod_{\ell|d} \Omega_\ell$$

for $d \geq 1$ squarefree, and for an arbitrary integrable function $x \mapsto \alpha(x)$, write

$$S_d(X; \alpha) = \int_{\{\rho_d(F_x) \in \Omega_d\}} \alpha(x) d\mu(x).$$

For $x \in X$, let $n(x) \geq 1$ be the integer defined by

$$n(x) = \prod_{\substack{\ell < L \\ \rho_d(F_x) \in \Omega_\ell}} \ell.$$

Notice that, for squarefree $d \in \mathcal{L}$, we have $\rho_d(F_x) \in \Omega_d$ if and only if $d \mid n(x)$.

Let

$$a_n = \int_{\{n(x)=n\}} \alpha(x) d\mu(x),$$

and then note the relation

$$S_d(X; \alpha) = \sum_{n \equiv 0 \pmod{d}} a_n.$$

Finally, define

$$P(L) = \prod_{\substack{\ell < L \\ \ell \in \mathcal{L}^*}} \ell.$$

Then we have

$$\begin{aligned} \int_{S(X, \Omega; L)} \alpha(x) d\mu(x) &= \int_{\{(n(x), P(L))=1\}} \alpha(x) d\mu(x) \\ &= \sum_{(n, P(L))=1} \left(\int_{\{n(x)=n\}} \alpha(x) d\mu(x) \right) = \sum_{(n, P(L))=1} a_n. \end{aligned}$$

Now let (λ_d^\pm) be two sequences of real numbers such that $\lambda_1^\pm = 1$ and

$$\sum_{d \mid n} \lambda_d^- \leq 0 \leq \sum_{d \mid n} \lambda_d^+$$

for $n \geq 2$. (The λ^+ are called upper-bound sieve coefficients, and the λ^- are called lower-bound sieve coefficients.) Then, if $\alpha(x) \geq 0$ for all x , we have

$$\sum_{(n, P(L))=1} a_n \leq \sum_n \left(\sum_{d \mid (n, P(L))} \lambda_d^+ \right) a_n = \sum_{d \mid P(L)} \lambda_d^+ \left(\sum_{n \equiv 0 \pmod{d}} a_n \right) = \sum_{d \mid P(L)} \lambda_d^+ S_d(X; \alpha),$$

and similarly

$$\sum_{(n, P(L))=1} a_n \geq \sum_{d \mid P(L)} \lambda_d^- S_d(X; \alpha).$$

It is natural to introduce the approximations (compare (2.10))

$$S_d(X; \alpha) = v_d(\Omega_d)H + r_d(X; \alpha) \tag{A.1}$$

(where v_d is the density as in Chapter 2), which is really a definition of $r_d(X; \alpha)$, where the ‘expected main term’ is

$$H = \int_X \alpha(x) d\mu(x).$$

Then, in effect, we have proved:

Proposition A.1 *Assume $\alpha(x) \geq 0$ for all $x \in X$. Let (λ_d^\pm) be arbitrary upper and lower-bound sieve coefficients which vanish for $d \geq D$, for some other parameter D . We have then*

$$V^-(\Omega)H - R^-(X; D) \leq \int_{S(X, \Omega; L)} \alpha(x) d\mu(x) \leq V^+(\Omega)H + R^+(X; D)$$

where

$$V^\pm(\Omega) = \sum_{d|P(L)} \lambda_d^\pm v_d(\Omega_d) \quad \text{and} \quad R^\pm(X; L) = \sum_{\substack{d < D \\ d|P(L)}} |\lambda_d^\pm r_d(X; \alpha)|.$$

But this is not quite what is needed for applications, because the main terms $V^\pm(X)$ are not yet in a form that makes them easy to evaluate. This next crucial step (usually called a ‘fundamental lemma’ in classical sieve theory) depends on the choice of λ_d^\pm (which is by no means obvious) and on properties of Ω_d . For instance, we have the following (see, e.g. [67, Corollary 6.2]; note this by no means the most general or best result known).

Proposition A.2 *Let $\kappa > 0$ and $y > 1$. There exist upper and lower-bound sieve coefficients (λ_d^\pm) , depending only on κ and y , supported on squarefree integers $< y$, bounded by one in absolute value, with the following properties: for all $s \geq 9\kappa + 1$ and $L^{9\kappa+1} < y$, we have*

$$\begin{aligned} \int_{S(X, \Omega; L)} \alpha(x) d\mu(x) &< \left(1 + e^{9\kappa+1-s} K^{10}\right) \prod_{\ell < L} (1 - v_\ell(\Omega_\ell)) H + R^+(X; L^s), \\ \int_{S(X, \Omega; L)} \alpha(x) d\mu(x) &> \left(1 - e^{9\kappa+1-s} K^{10}\right) \prod_{\ell < L} (1 - v_\ell(\Omega_\ell)) H + R^-(X; L^s), \end{aligned}$$

provided the sieving sets (Ω_ℓ) satisfy the condition

$$\prod_{w \leq \ell < L} (1 - v_\ell(\Omega_\ell))^{-1} \leq K \left(\frac{\log L}{\log w} \right)^\kappa, \quad \text{for all } w \text{ and } L, 2 \leq w < L < y, \quad (\text{A.2})$$

for some $K \geq 0$.

In standard applications, $r_d(X; \alpha)$ should be ‘small’,¹ as the remainder term in some equidistribution theorem. Note again that this can only be true if the family (ρ_d) is linearly disjoint. If this remainder is well-controlled on average over $d < D$, for some D (as large as possible) we can apply the above for L such that $L^s < D$ (with $s \geq 9\kappa + 1$). Note that when s is large enough (i.e., L small enough), the coefficient $1 \pm e^{9\kappa+1-s} K^{10}$ will be close to 1, in particular it will be *positive* in the lower bound.

Further, the condition (A.2) holds if $v_\ell(\Omega_\ell)$ is of size $\kappa \ell^{-1}$ on average. This is the traditional context of a small sieve of dimension κ ; we see that in the abstract framework, this means rather that the sieving sets Ω_ℓ are ‘of codimension 1’ in a certain sense. The important case $\kappa = 1$ (the classical ‘linear sieve’) corresponds intuitively to sieving sets defined by a single irreducible algebraic condition.

We recall (see Section 2.5) that the factor

$$\prod_{\ell < L} (1 - v_\ell(\Omega_\ell))$$

is the natural one to expect intuitively if $v_\ell(\Omega_\ell)$ is interpreted as the probability of $\rho_\ell(F_x)$ being in Ω_ℓ , and if the various ℓ are independent. Recall also that if \mathcal{L} is the full power set of the prime sieve support \mathcal{L}^* , then the saving factor H in (2.4) is given by

$$H^{-1} = \prod_{\ell \in \mathcal{L}^*} (1 - v_\ell(\Omega_\ell)).$$

Finally, some words concerning the Selberg sieve. We do not give details, since there are many excellent presentations in the literature, and readers would have no trouble adapting them to the general sieve setting, using all the previous work. A few points deserve mention: first, just as in the classical case, the Selberg sieve is a priori an upper-bound sieve, and one needs to use some type of Buchstab identity to transform it to a lower-bound sieve; second, just as the large sieve can be used as upper-bound sieve even in small sieve contexts, so is the Selberg sieve applicable in large sieve contexts. In fact, much of the qualitative part of the theory of Chapter 2 and of its applications in this book could have been developed using a general Selberg sieve. The exception is the dual sieve which (to the author’s knowledge) is really a feature of the large sieve. Also, the qualitative similarity does not extend to the finest quantitative results. Indeed, the Selberg sieve starts from assumptions such as (A.1), which are akin

¹ Possibly only on average over d , since this is how those terms occur in the sieve remainder. This is a crucial feature, for instance, in the study of primes in arithmetic progressions, where the Bombieri–Vinogradov Theorem leads to estimates which are on average as strong as the Generalized Riemann Hypothesis allows (and even stronger results are known, due to Fouvry–Iwaniec, Bombieri–Friedlander–Iwaniec).

to the individual equidistribution assumptions of Section 2.3. In many deep applications, those statements are in fact the most crucial part, and they are (or will most likely be) proved by applying the Weyl criterion for equidistribution, hence, by estimating suitable ‘exponential sums’ similar to the $W(\varphi, \varphi')$. In applications like those in Chapters 7 and 8, any impression of greater simplicity in using one sieve or another seems to be a minor issue compared with the depth of the tools involved.

A.2 An application

To illustrate the use of lower-bound sieves, we conclude with a simple application related to Theorems 8.15 and 8.22 in Chapter 8. The reader will have no problem supplying a similar result in the context of sieve for random walks on $SL(n, \mathbf{Z})$ or $Sp(2g, \mathbf{Z})$.

Proposition A.3 *Let q be a power of a prime number $p \geq 5$, $g \geq 1$ an integer and let $f \in \mathbf{F}_q[T]$ be a squarefree polynomial of degree $2g$. For t not a zero of f , let C_t denote the smooth projective model of the hyperelliptic curve $y^2 = f(x)(x - t)$, and let J_t denote its Jacobian variety. There exists an absolute constant $\alpha \geq 0$ such that*

$$|\{u \in \mathbf{F}_q \mid f(t) \neq 0 \text{ and } |C_t(\mathbf{F}_q)| \text{ has no odd prime factor } < q^\gamma\}| \gg \frac{q}{\log q},$$

$$|\{u \in \mathbf{F}_q \mid f(t) \neq 0 \text{ and } |J_t(\mathbf{F}_q)| \text{ has no odd prime factor } < q^\gamma\}| \gg \frac{q}{\log q}$$

for any γ such that

$$\gamma^{-1} > \alpha(2g^2 + g + 1)(\log \log 3g),$$

where the implied constants depend only on g and γ .

In particular, for any fixed g , there are infinitely many points $t \in \bar{\mathbf{F}}_q$ such that $|C_t(\mathbf{F}_{q^{\deg(t)}})|$ has at most $\alpha(2g^2 + g + 1)(\log \log 3g) + 2$ prime factors, and similarly for $|J_t(\mathbf{F}_{q^{\deg(t)}})|$.

Remark A.4

- (1) It may well be that $|J_t(\mathbf{F}_q)|$ is even for all t , since if f has a root x_0 in \mathbf{F}_q , it will define a non-zero point of order 2 in $J_t(\mathbf{F}_q)$.
- (2) There are results, due to Cojocaru [24] in particular, giving almost prime values of group orders of reductions of elliptic curves over \mathbf{Q} ; except for curves with complex multiplication, they are currently conditional on the Generalized Riemann Hypothesis.

Proof Obviously, we wish to use the same coset sieve setting and siftable set as in Theorems 8.15 and Theorem 8.22, and consider the sieving sets

$$\Omega_\ell^J = \{g \in CSp(2g, \mathbf{F}_\ell) \mid g \text{ is } q\text{-symplectic and } \det(g - 1) = 0 \in \mathbf{F}_\ell\},$$

$$\Omega_\ell^C = \{g \in CSp(2g, \mathbf{F}_\ell) \mid g \text{ is } q\text{-symplectic and } \mathrm{Tr}(g) = q + 1\},$$

for $\ell \geq 3$. By (5) and (6) of Proposition B.4, we have

$$v_\ell(\Omega_\ell^{\mathbf{S}}) = \frac{|\Omega_\ell^{\mathbf{S}}|}{|Sp(2g, \mathbf{F}_\ell)|} \leq \min\left(1, \frac{1}{\ell} \left(\frac{\ell}{\ell-1}\right)^{4g^2}\right),$$

where $\mathbf{S} \in \{C, J\}$, but since the stronger upper bound only becomes effective for ℓ large enough, we replace Ω_ℓ by the empty set for small ℓ . Precisely, it is not difficult to check that there exists an absolute constant $A > 0$ such that if $L_0 = Ag^2 \log 2g$, we have

$$\begin{aligned} \prod_{L_0 < \ell \leq L} \left(1 - \frac{1}{\ell} \left(\frac{\ell}{\ell-1}\right)^{4g^2}\right)^{-1} &\ll \prod_{L_0 < \ell \leq L} \left(1 - \frac{1}{\ell}\right)^{-1} \prod_{\ell > L_0} \left(1 - \frac{g^2}{\ell^2}\right) \\ &\ll \log L \end{aligned}$$

for all $g \geq 2$ and $L \geq L_0$, with an absolute implied constant.

We take $\Omega_\ell^{\mathbf{S}} = \emptyset$ for all $\ell < L_0$, and keep the previous ones for $\ell \geq L_0$. Then it is easily checked that (A.2) holds with $\kappa = 1$ and $K \ll 1$ (consider separately $L < L_0$ and $L \geq L_0$ and use the preceding estimate).

Coming to the error term $R^-(X; L)$, individual estimates for $r_d(X; \alpha)$ with $\alpha(x) = 1$ amount to estimates for the error term in the Chebotarev density theorem (which is the individual equidistribution in this context, as in Remark 2.14). Using Proposition 8.8 (see also Theorem 1.3 in [81]), we obtain

$$r_d(X; \alpha) \ll gq^{1/2} |\Omega_d^{\mathbf{S}}|^{1/2} \ll gq^{1/2} \left(\psi(d)^{2g^2+g} d^{g-1} \varphi(d)^{-g} \right)^{1/2},$$

with absolute implied constants, and hence

$$R^-(X; L^s) \ll gq^{1/2} L^{s(2g^2+g+1)/2} (\log \log 3L^s)^{g^2+g},$$

for any $s \geq 1$, with an absolute implied constant.

Let $s = \log 2 + 10 \log K \ll \log \log 3g$, and let $\varepsilon > 0$ be arbitrarily small. Then we can take

$$L = q^{(s(2g^2+g+1))^{-1-\varepsilon}}$$

in the lower-bound sieve, which gives

$$|\{t \in \mathbf{F}_q \mid f(t) \neq 0 \text{ and } |\mathbf{S}_t(\mathbf{F}_q)| \text{ has no odd prime factor } < L\}| \\ \gg q \prod_{\ell < L} (1 - v_\ell(\Omega_\ell^S)) \gg q \prod_{L_0 < \ell < L} \left(1 - \frac{1}{\ell} \left(\frac{\ell}{\ell+1}\right)^{4g^2}\right) \gg \frac{q}{\log q}$$

provided $L > L_0 = Ag^2 \log 2g$ and with absolute implied constants. Putting all together, the theorem follows easily. \square

Appendix B

Local density computations over finite fields

B.1 Density of cycle types for polynomials over finite fields

We recall the basic counting lemma for polynomials over a finite field with a given factorization type, giving the uniform version proved in [80, Lemma 7.3 (i)] (with some refinements).

Lemma B.1 *Let ℓ be a prime number, $r \geq 1$ an integer. Let $n_i \geq 0$, $1 \leq i \leq r$, be integers such that*

$$r = n_1 + 2n_2 + \cdots + rn_r.$$

The cardinality of the set Ω_ℓ of monic polynomials $f \in \mathbf{F}_\ell[T]$ which factor as a product

$$f = f_1 \cdots f_r,$$

where f_i , $1 \leq i \leq r$, is a product of n_i distinct irreducible monic polynomials of degree i , satisfies

$$\frac{|c|}{|\mathfrak{S}_r|} \ell^r \left(1 - \frac{1}{\ell}\right)^{n_2 + \sum n_i} \left(1 - \frac{1}{\sqrt{\ell}}\right)^{n_1} \leq |\Omega_\ell| \leq \frac{|c|}{|\mathfrak{S}_r|} \ell^r, \quad (\text{B.1})$$

for all $\ell > r^2$, and for $\ell > 4r$ if $n_1 = 0$, where c is the conjugacy class in \mathfrak{S}_r of permutations whose expression as a product of disjoint cycles involves n_i cycles of length i , $1 \leq i \leq r$, precisely

$$|c| = r! \prod_{1 \leq i \leq r} \frac{1}{i^{n_i} n_i!}.$$

In particular, as $\ell \rightarrow +\infty$, we have

$$|\Omega_\ell| \sim \frac{|c|}{\mathfrak{S}_r} \ell^r.$$

For the counting of irreducible polynomials, we have $n_1 = \cdots = n_{r-1} = 0$, $n_r = 1$, and

$$\frac{1}{r} \ell^r \left(1 - \frac{1}{\ell}\right) \leq |\{f \in \mathbf{F}_\ell[T] \mid f \text{ is irreducible, monic, of degree } r\}| \leq \frac{1}{r} \ell^r,$$

where the lower bound holds for all $\ell > 4r$.

Proof By unique factorization in $\mathbf{F}_\ell[T]$, we have

$$|\Omega_\ell| = \prod_{1 \leq i \leq r} \binom{p(i, \ell)}{n_i},$$

where $p(i, \ell)$ is the number of irreducible monic polynomials of degree i in $\mathbf{F}_\ell[T]$. This latter quantity is expressed by a classical formula of Gauss:

$$p(i, \ell) = \frac{1}{i} \sum_{d|i} \mu(d) \ell^{i/d};$$

this is the expression, by inclusion-exclusion (or Möbius inversion), of the partition of the extension \mathbf{F}_{ℓ^i} of \mathbf{F}_ℓ by means of elements which generate the subextensions of degree d , $d \mid i$, and of the fact that $ip(i, \ell)$ is precisely the number of such elements since they are themselves partitioned into $p(i, \ell)$ sets of i roots of irreducible polynomials of degree i .

Using this, it is clear that $p(i, \ell) \leq \frac{1}{i} \ell^i$, and so we have

$$\binom{p(i, \ell)}{n_i} \leq \frac{1}{n_i!} \left(\frac{\ell^i}{i}\right)^{n_i},$$

from which the upper bound in (B.1) follows (without any condition on the size of ℓ compared with r).

For the lower bound, we claim the following:

$$p(1, \ell) \geq \ell \left(1 - \frac{1}{\sqrt{\ell}}\right) + r - 1, \quad \text{for } \ell > 4r^2, \quad (\text{B.2})$$

$$p(2, \ell) \geq \frac{\ell^2}{2} \left(1 - \frac{1}{\ell}\right)^2 + \frac{r}{2} - 1, \quad \text{for } \ell > r, \quad (\text{B.3})$$

$$p(i, \ell) \geq \frac{\ell^i}{i} \left(1 - \frac{1}{\ell}\right) + \frac{r}{i} - 1, \quad \text{for } 3 \leq i \leq r, \ell > 4r. \quad (\text{B.4})$$

To see this, we consider $i = 1$ and $i = 2$ separately, namely

$$p(1, \ell) = \ell \geq \ell \left(1 - \frac{1}{\sqrt{\ell}}\right) + r - 1$$

for $\ell > r^2$ (by inspection), and

$$p(2, \ell) = \frac{1}{2}(\ell^2 - \ell) \geq \frac{1}{2}(\ell - 1)^2 + \frac{r}{2} - 1$$

if $\ell \geq r$. For $i \geq 3$, we have

$$p(i, \ell) = \frac{1}{i} \ell^i + \frac{1}{i} \sum_{\substack{d|i \\ d < i}} \mu(d) \ell^{d/i} \geq \frac{1}{i} \ell^i - \frac{\sqrt{\ell}}{i} \sum_{d < i} 1 \geq \frac{1}{i} \ell^i - \sqrt{\ell},$$

by the Gauss formula. Hence it suffices to show that

$$\frac{\ell^i}{i} - \ell^{i/2} > \frac{\ell^i}{i} \left(1 - \frac{1}{\ell}\right) + \frac{r}{i}$$

for $\ell > 4r$ in order to obtain (B.4). This amounts to

$$\frac{\ell^{i-1}}{i} > \ell^{i/2} + \frac{r}{i},$$

which we check as follows for $i \geq 4$,

$$\frac{\ell^{i-1}}{i} > 2\ell^{i-2} \geq \ell^{i/2} + \ell^{i-2} \geq \ell^{i/2} + \ell \geq \ell^{i/2} + \frac{r}{i},$$

leaving the case $i = 3$ to the reader.

From (B.2) and (B.4), we now derive for all i , $3 \leq i \leq r$ and all $n_i \leq r/i$ that

$$\begin{aligned} \binom{p(i, \ell)}{n_i} &= \frac{p(i, \ell)(p(i, \ell) - 1) \cdots (p(i, \ell) - n_i + 1)}{n_i!} \\ &\geq \frac{(p(i, \ell) - r/i + 1)^{n_i}}{n_i!} \geq \left(1 - \frac{1}{\ell}\right)^{n_i} \frac{1}{i^{n_i} n_i!} \ell^{in_i} \end{aligned}$$

if $\ell > 4r$, and for $i = 1$, $n_1 \leq r$, that

$$\begin{aligned} \binom{p(1, \ell)}{n_1} &= \frac{p(1, \ell)(p(1, \ell) - 1) \cdots (p(1, \ell) - n_1 + 1)}{n_1!} \\ &\geq \frac{(p(1, \ell) - r + 1)^{n_1}}{n_1!} \geq \left(1 - \frac{1}{\sqrt{\ell}}\right)^{n_1} \frac{1}{1^{n_1} n_1!} \ell^{n_1} \end{aligned}$$

for $\ell > r^2$, and finally for $i = 2$, $n_2 \leq r/2$, $\ell \geq r$, that

$$\begin{aligned} \binom{p(2, \ell)}{n_2} &= \frac{p(2, \ell)(p(2, \ell) - 1) \cdots (p(2, \ell) - n_2 + 1)}{n_2!} \\ &\geq \frac{(p(2, \ell) - r/2 + 1)^{n_2}}{n_2!} \geq \left(1 - \frac{1}{\ell}\right)^{2n_2} \frac{1}{2^{n_2} n_2!} \ell^{2n_2}. \end{aligned}$$

Hence, putting these together, we get

$$\begin{aligned} |\Omega_\ell| &\geq \left(1 - \frac{1}{\sqrt{\ell}}\right)^{n_1} \left(1 - \frac{1}{\ell}\right)^{n_2 + \sum n_i} \left(\prod_{1 \leq i \leq r} \frac{1}{i^{n_i} n_i!}\right) \ell^{\sum i n_i} \\ &= \frac{|c|}{|\mathfrak{S}_r|} \left(1 - \frac{1}{\sqrt{\ell}}\right)^{n_1} \left(1 - \frac{1}{\ell}\right)^{n_2 + \sum n_i} \ell^r, \end{aligned}$$

under the stated conditions on ℓ . \square

Here is a result for polynomials with a fixed value at 0, which will be used when dealing with characteristic polynomials of unimodular matrices.

Lemma B.2 *Let ℓ be a prime number, $r \geq 1$ an integer. Let $n_i \geq 0$, $1 \leq i \leq r$, be integers such that*

$$r = n_1 + 2n_2 + \cdots + rn_r.$$

The cardinality of the set Ω_ℓ^1 of monic polynomials $f \in \mathbf{F}_\ell[T]$ such that $f(0) = 1$ and which factor as a product

$$f = f_1 \cdots f_r$$

where f_i , $1 \leq i \leq k$, is a product of n_i distinct irreducible monic polynomials of degree i , satisfies

$$|\Omega_\ell^1| \geq \frac{|c|}{|\mathfrak{S}_r|} \ell^{r-1} \left(1 - \frac{1}{\ell}\right)^{n_2 + \sum n_i + 1} \left(1 - \frac{1}{\sqrt{\ell}}\right)^{n_1}, \quad (\text{B.5})$$

for all $\ell > 16r^2$, where c is the conjugacy class in \mathfrak{S}_r of permutations whose expression as a product of disjoint cycles involves n_i cycles of length i , $1 \leq i \leq r$.

Proof The proof is very similar to that of Lemma B.1, but requires a fair number of small checks; the reader should at least check quickly that the asymptotic version of the inequality is quite obvious. The idea is that we can select (with few limitations) all but one of the irreducible factors as in Lemma B.1, and ensure that the condition that the constant coefficient is 1 holds by selecting the last factor among those with the right constant coefficient (which is not fixed, however).

To start with some notation, let $q(i, \ell, a)$, for $a \in \mathbf{F}_\ell^\times$, denote the number of irreducible monic polynomials of degree i in $\mathbf{F}_\ell[T]$ with constant coefficient a . We distinguish three cases: (i) $n_1 = 0$; (ii) $n_1 \neq 0$ and $n_2 = \cdots = n_r = 0$; (iii) $n_1 \neq 0$ and $n_i \neq 0$ for some i with $2 \leq i \leq r$. (The reason for considering n_1 separately is that X is the only irreducible polynomial with zero constant term, and so it can not occur as a factor in a product with constant term 1.)

We start with the first case. Let $s \geq 2$ be the largest integer such that $n_s \neq 0$. Then we have the following lower bound

$$|\Omega_\ell^1| \geq \prod_{2 \leq i \leq s-1} \binom{p(i, \ell)}{n_i} \times \binom{p(s, \ell)}{n_s - 1} \times \min_{a \in \mathbf{F}_\ell^\times} \frac{q(s, \ell, a) - n_s + 1}{n_s}, \quad (\text{B.6})$$

which corresponds to the previous idea. Note that as before we have $n_s \leq r/s$.

We now give lower bounds for $q(s, \ell, a)$ for $s \geq 2$. This number is also equal to the number of Galois orbits of elements of norm a in \mathbf{F}_{ℓ^s} with are of degree s and not smaller. Since the norm map $\mathbf{F}_{\ell^s}^\times \rightarrow \mathbf{F}_\ell^\times$ is onto, we see that we have the following lower bounds:

$$q(s, \ell, a) \geq \frac{1}{s} \frac{\ell^s - 1}{\ell - 1} - \ell^{s/2} \quad (\text{B.7})$$

$$\geq \frac{\ell^{s-1}}{s} - \ell^{s/2}. \quad (\text{B.8})$$

We now claim that we have

$$q(s, \ell, a) \geq \frac{1}{\ell} \left(1 - \frac{1}{\ell}\right) \frac{\ell^s}{s} + \frac{r}{s} - 1, \quad \text{if } s \geq 3, \ell > 2r \quad (\text{B.9})$$

$$q(2, \ell, a) \geq \frac{\ell}{2} \left(1 - \frac{1}{\sqrt{\ell}}\right), \quad \text{if } \ell > 16r^2. \quad (\text{B.10})$$

This, together with (B.6), proves the lemma in the case $n_1 = 0$. Now, using (B.8), we can check (B.9) for $s \geq 5$ in the same manner that (B.4) was checked. For $s = 4$, we can use the refinement

$$q(4, \ell, a) \geq \frac{1}{4} \frac{\ell^4 - 1}{\ell - 1} - \frac{\ell^2}{2}$$

of (B.7), and proceed directly. For $s = 3$, we may notice that elements in $\mathbf{F}_{\ell^3}^\times$ are either of degree 1 or 3, and there are at most three elements in \mathbf{F}_ℓ^\times with norm a (from $\mathbf{F}_{\ell^3}^\times$), so that

$$q(3, \ell, a) \geq \frac{\ell^2}{3} - 3,$$

which is again sufficient to obtain (B.9). Finally, for $s = 2$, there are at most two exceptional elements, and thus $q(2, \ell, a) \geq \frac{1}{2}(\ell + 1) - 2$, which gives (B.10) straight away.

This concludes the analysis of the first case; we pass to the second case, where $n_2 = \dots = n_r = 0$, i.e., we are counting polynomials which are products of $n_1 = r$ linear factors. Then we claim that

$$|\Omega_\ell^1| \geq \frac{(\ell - 1) \cdots (\ell - 1 - (r - 2) + 1)}{r!} \times (\ell - 1 - 2r + 2). \quad (\text{B.11})$$

Indeed, we can choose the first factor $X + \alpha_1$ among the $\ell - 1$ polynomials $X + \alpha$ with $\alpha \in \mathbb{F}_\ell^\times$, the second $X + \alpha_2$ among the $\ell - 2$ remaining such polynomials, and so on, up to the $(r - 2)$ -th factor $X + \alpha_{r-2}$. However, the $(r - 1)$ -th factor $X + \beta$ is subject to some constraints, beyond being different from the $r - 2$ previous ones, because after it is selected, the last factor is necessarily $X + \gamma$ with

$$\alpha_1 \cdots \alpha_{r-2} \beta \gamma = 1.$$

This fixes γ , except that the choice may be forbidden, if either $\gamma = \alpha_i$, $1 \leq i \leq r - 2$, or $\gamma = \beta$. To avoid the first case means that

$$\beta \neq \frac{1}{\alpha_i} \frac{1}{\alpha_1 \cdots \alpha_{r-2}}$$

and to avoid the second, that

$$\beta^2 \neq \frac{1}{\alpha_1 \cdots \alpha_{r-2}};$$

hence there are $\leq 2 + r - 2 = r$ additional possible excluded factors.

Now since $\ell - 1 - 2r + 1 > \ell(1 - \ell^{-1/2})$ for $\ell > 4r^2$, we deduce that

$$|\Omega_\ell^1| \geq \frac{\ell^{r-1}}{r!} \left(1 - \frac{1}{\sqrt{\ell}}\right)^r$$

for $\ell > 4r^2$.

Finally we conclude with the last case. Then, $s \geq 2$ being defined as before, the bound (B.6) may be replaced with

$$|\Omega_\ell^1| \geq \binom{p(1, \ell) - 1}{n_1} \times \prod_{2 \leq i \leq s-1} \binom{p(i, \ell)}{n_i} \times \binom{p(s, \ell)}{n_s - 1} \times \min_{a \in \mathbb{F}_\ell^\times} \frac{q(s, \ell, a) - n_s + 1}{n_s},$$

since we need only make sure of not using X as a linear factor. Since

$$\binom{p(1, \ell) - 1}{n_1}$$

is at least as large as the lower bound (B.11), we can conclude by combining the two previous cases. \square

Remark B.3 More generally (but only asymptotically for a fixed r), Rivin [108] has shown that for any j , $0 \leq j \leq r - 1$, among monic polynomials of degree r with a given factorization type, the proportion of those with a fixed j -th coefficient is equivalent with ℓ^{-1} as ℓ tends to infinity.

B.2 Some matrix densities over finite fields

In Chapters 7, 8, and in Appendix A, we have quoted various estimates for the ‘density’ of certain subsets of matrix groups over finite fields, which are required to prove lower (or upper) bounds for the saving factor H in certain applications of the large sieve inequalities. We prove those statements here, relying mostly on the work of Chavdarov [22] to link such densities with those of polynomials of certain types, which are much easier to compute. In one case, however, we use the Riemann Hypothesis over finite fields to estimate a multiplicative exponential sum.

Proposition B.4 *Let $\ell \geq 2$ be a prime number.*

- (1) *Let $G = SL(n, \mathbf{F}_\ell)$ or $G = Sp(2g, \mathbf{F}_\ell)$, with $n \geq 2$ or $g \geq 1$. Then we have*

$$\frac{1}{|G|} |\{x \in G \mid \det(T - x) \in \mathbf{F}_\ell[T] \text{ is irreducible}\}| \gg 1$$

where the implied constant depends only on n or g , and more precisely for $G = SL(n, \mathbf{F}_\ell)$ we have

$$\frac{1}{|G|} |\{x \in G \mid \det(T - x) \in \mathbf{F}_\ell[T] \text{ is irreducible}\}| \geq \frac{1}{n} \left(\frac{\ell}{\ell + 1} \right)^{n^2} \left(1 - \frac{1}{\ell} \right)$$

for all $\ell > n$.

- (2) *Let $G = SL(n, \mathbf{F}_\ell)$ or $G = Sp(2g, \mathbf{F}_\ell)$, with $n \geq 2$ or $g \geq 1$; let i, j be integers with $1 \leq i, j \leq n$ or $1 \leq i, j \leq 2g$, respectively. Then, if $\ell \geq 3$, we have*

$$\frac{1}{|G|} |\{x = (x_{\alpha, \beta}) \in G \mid x_{i, j} \in \mathbf{F}_\ell \text{ is not a square}\}| \gg 1$$

where the implied constant depends only on n or g . In the next results (3), (4), (5), (6), let $G = CSp(2g, \mathbf{F}_\ell)$ with $g \geq 1$, $d = 2g^2 + g + 1$ its dimension, and denote by $m(x) \in \mathbf{F}_\ell^\times$ the multiplier of a symplectic similitude $x \in G$.

- (3) *For any $q \in \mathbf{F}_\ell^\times$, we have*

$$\frac{1}{|Sp(2g, \mathbf{F}_\ell)|} |\{x \in G \mid m(x) = q, \det(x - 1) \text{ is a square in } \mathbf{F}_\ell\}| \geq \frac{1}{2} \left(\frac{\ell}{\ell + 1} \right)^d.$$

- (4) *For any $q \in \mathbf{F}_\ell^\times$, we have*

$$\frac{1}{|Sp(2g, \mathbf{F}_\ell)|} |\{x \in G \mid m(x) = q, q + 1 - \text{Tr}(x) \text{ is a square in } \mathbf{F}_\ell\}| \geq \frac{1}{2} \left(\frac{\ell}{\ell + 1} \right)^d.$$

- (5) *For any $q \in \mathbf{F}_\ell^\times$, we have*

$$\frac{1}{|Sp(2g, \mathbf{F}_\ell)|} |\{x \in G \mid m(x) = q, \det(x - 1) = 0\}| \leq \min \left(1, \frac{1}{\ell} \left(\frac{\ell}{\ell - 1} \right)^d \right).$$

(6) For any $q \in \mathbf{F}_\ell^\times$, we have

$$\frac{1}{|Sp(2g, \mathbf{F}_\ell)|} |\{x \in G \mid m(x) = q, q + 1 - \text{Tr}(x) = 0\}| \leq \min\left(1, \frac{1}{\ell} \left(\frac{\ell}{\ell-1}\right)^d\right).$$

(7) Finally, for any fixed Lagrangian subspace $J \subset \mathbf{F}_\ell^{2g}$, we have

$$\frac{1}{|Sp(2g, \mathbf{F}_\ell)|} |\{x \in Sp(2g, \mathbf{F}_\ell) \mid xJ \text{ is transverse to } J\}| = \prod_{j=1}^g \frac{1}{1 + \ell^{-j}}.$$

For all except the second and last points, the following result due to Chavdarov [22, Section 3] is the crucial point: it expresses in a very precise manner the fact that characteristic polynomials of matrices in the finite groups of Lie type we consider are equidistributed among the ‘obvious’ candidate polynomials. Recall that a semisimple matrix is simply a diagonalizable matrix.

Lemma B.5 *Let $\mathbf{G} = GL(n)$ or $G = CSp(2g)$ over \mathbf{F}_ℓ , $r = \text{rank } \mathbf{G}$, which is n , or $g + 1$, $d = \dim \mathbf{G}$, which is n^2 or $2g^2 + g + 1$, respectively. Let $\mathbf{G}' = SL(n)$ or $Sp(2g)$ be the derived group.*

(1) *Let $f_0 \in \mathbf{F}_\ell[T]$ be the characteristic polynomial of a semisimple element $g_0 \in \mathbf{G}(\mathbf{F}_\ell)$. Then we have*

$$\frac{|\mathbf{G}'(\mathbf{F}_\ell)|}{\ell^{r-1}} \left(\frac{\ell}{\ell+1}\right)^d \leq |\{x \in \mathbf{G}(\mathbf{F}_\ell) \mid \det(T - x) = f_0\}| \leq \frac{|\mathbf{G}'(\mathbf{F}_\ell)|}{\ell^{r-1}} \left(\frac{\ell}{\ell-1}\right)^d.$$

(2) *Let $X_{\mathbf{G}}$ be the following subsets of polynomials in $\mathbf{F}_\ell[T]$:*

$$X_{GL(n)} = \{f \mid \deg(f) = n, f \text{ monic}\},$$

$$X_{CSp(2g)} = \{f \mid \deg(f) = 2g, f \text{ symplectic}\},$$

where¹ a monic polynomial f is symplectic if it satisfies

$$f\left(\frac{q}{T}\right) = q^s T^{-2s} f(T)$$

for some invertible element q . Then for any $f_0 \in X_{\mathbf{G}}$, there exists a semisimple element $g_0 \in \mathbf{G}(\mathbf{F}_\ell)$ such that $\det(T - g_0) = f_0$.

Note that Lemma 7.2 in [80], which is the analogue of the lower bound in (1), is in error (it misses the factor $(\frac{\ell}{\ell+1})^d$, essentially; this does not seriously affect the paper . . .).

¹ In terms of the reversed characteristic polynomial, as discussed in Chapter 8 (see (8.3)), this means that $T^{-2s} f(T^{-1})$ is q -symplectic in the sense of (8.3), for some q .

Proof

(1) This is essentially an application of the method of the proof of Theorem 3.5 in [22], which is attributed to Borel and which corresponds to the case of $CSp(2g, \mathbf{F}_\ell)$. We indicate the strategy of the proof for $G = GL(n, \mathbf{F}_\ell)$, or indeed for $SL(n, \mathbf{F}_\ell)$ in view of the argument used. Let $q = \det(g_0)$, and let Δ be the set of those $g \in GL(n, \mathbf{F}_\ell)$ with characteristic polynomial equal to f_0 ; note that $\det(g) = \det(g_0) = q$ for all $g \in \Delta$. The elements of Δ can be parametrized by pairs (g_s, g_u) where:

- g_s is a semisimple element of G which is $SL(n, \mathbf{F}_\ell)$ -conjugate to g_0 ;
- g_u is a unipotent element (i.e., for some N , we have $(g_u - \text{Id})^N = \text{Id}$) of G which commutes with g_s .

Indeed, we can map such pairs to Δ by taking $g = g_s g_u$, which has the same characteristic polynomial as g_s , hence as g_0 ; and conversely, any g can be expressed uniquely by Jordan decomposition as a product $g = g_s g_u$, with g_s semisimple and g_u unipotent, which commute. Then the equality $f_0 = \det(T - g) = \det(T - g_s)$ implies that g_0 and g_s are conjugate by an element of $SL(n, \mathbf{F}_\ell)$ (not over the algebraic closure).

Now we count the pairs (g_s, g_u) . For each given g_s , the possibilities for g_u are all the unipotent elements in the group of \mathbf{F}_ℓ -points of the centralizer

$$C(g_s) = \{h \in GL(n, \overline{\mathbf{F}}_\ell) \mid hg_s = g_sh\}$$

of g_s in $GL(n, \overline{\mathbf{F}}_\ell)$. A theorem of Steinberg, depending crucially on the fact that $SL(n)$ is simply-connected, states that this centralizer is connected (as algebraic group) and contains precisely $\ell^{\delta-n}$ unipotent elements, where δ is the dimension of the centralizer, and n is the rank of $C(g_s)$ (the rank is n because the centralizer contains a maximal torus of rank n , namely any maximal torus in $GL(n)$ containing g_s). The dimension does not depend on g_s since by assumption all g_s are conjugate, hence have isomorphic centralizers.

On the other hand, the number of possibilities for g_s is the order of the orbit of g_0 under $SL(n, \mathbf{F}_\ell)$ conjugation, which is

$$\frac{|SL(n, \mathbf{F}_\ell)|}{|C_0(\mathbf{F}_\ell)|}$$

where

$$C_0 = \{h \in SL(n, \overline{\mathbf{F}}_\ell) \mid hg_0h^{-1} = g_0\}.$$

Thus we find

$$|\Delta| = \ell^{d-n} \frac{|SL(n, \mathbf{F}_\ell)|}{|C_0(\mathbf{F}_\ell)|},$$

and it now suffices to observe that C_0 is, up to scalars, isomorphic to any $C(g_s)$; this shows that C_0 is of dimension $\delta - 1$, and a result of Serre (using once more the connectedness of C_0) gives

$$(\ell - 1)^{\delta-1} \leq |C_0(\mathbf{F}_\ell)| \leq (\ell + 1)^{\delta-1}.$$

Using the fact that $\delta \leq d$, we obtain the result as stated.

- (2) This is proved by Chaydarov in [22, Lemmas 3.4, 3.8] (though the results are stated for the reversed characteristic polynomials, which is of course irrelevant). Note that if f is assumed to have distinct roots,² which is the generic situation, the proof is simpler since any matrix with coefficients in \mathbf{F}_ℓ and characteristic polynomial f will work, e.g., the companion matrix for $GL(n)$. \square

Proof of Proposition B.4 (1) Take the case $G = SL(n, \mathbf{F}_\ell)$, for instance (for $Sp(2g, \mathbf{F}_\ell)$, since we do not state a uniform estimate with respect to g , the result is much simpler). We need to compute

$$\frac{1}{|G|} \sum_{f \in \tilde{\Omega}_\ell} |\{g \in G \mid \det(T - g) = f\}|,$$

where f runs over the set $\tilde{\Omega}_\ell$ of irreducible monic polynomials $f \in \mathbf{F}_\ell[T]$ of degree n with $f(0) = 1$. By the previous lemma, using the fact that all elements in $GL(n, \mathbf{F}_\ell)$ with characteristic polynomial in $\tilde{\Omega}_\ell$ are in $SL(n, \mathbf{F}_\ell)$, we have

$$|\{g \in G \mid \det(T - g) = f\}| \geq \frac{|G|}{\ell^{n-1}} \left(\frac{\ell}{\ell + 1}\right)^{n^2}$$

for all $f \in \tilde{\Omega}_\ell$, and by Lemma B.2, we obtain

$$\begin{aligned} \frac{1}{|G|} \sum_{f \in \tilde{\Omega}_\ell} |\{g \in G \mid \det(T - g) = f\}| &\geq \frac{1}{\ell^{n-1}} \left(\frac{\ell}{\ell + 1}\right)^{n^2} |\tilde{\Omega}_\ell| \\ &\geq \frac{1}{n} \left(\frac{\ell}{\ell + 1}\right)^{n^2} \left(1 - \frac{1}{\ell}\right) \end{aligned}$$

for all $\ell > n$.

- (2) By detecting squares using the Legendre character, we need to compute

$$\frac{1}{2|G|} \sum_{\substack{g \in G \\ g_{i,j} \neq 0}} \left(1 + \left(\frac{g_{i,j}}{\ell}\right)\right)$$

² This is the case of interest when looking at matrices with irreducible characteristic polynomials.

where (\cdot) is the non-trivial quadratic character of \mathbf{F}_ℓ^\times . Let \mathbf{G} be the algebraic group $SL(n)$ or $Sp(2g)$ over \mathbf{F}_ℓ , d its dimension (either $n^2 - 1$ or $2g^2 - g$). Since $\mathbf{G} \cap \{g_{i,j} = 0\}$ is obviously a proper closed subset of the geometrically connected affine variety \mathbf{G} , the affine variety

$$\mathbf{G}_{i,j} = \mathbf{G} - \mathbf{G} \cap \{g_{i,j} = 0\}$$

over \mathbf{F}_ℓ is geometrically connected of dimension d , and we have

$$|\mathbf{G}_{i,j}(\mathbf{F}_\ell)| = |\{g \in \mathbf{G}(\mathbf{F}_\ell) \mid g_{i,j} \neq 0\}| \gg |\mathbf{G}(\mathbf{F}_\ell)|,$$

for $\ell \geq 3$. This means that it is enough to prove

$$\sum_{g \in \mathbf{G}_{i,j}(\mathbf{F}_\ell)} \left(\frac{g_{i,j}}{\ell} \right) \ll \ell^{d-1/2}$$

for $\ell \geq 3$, the implied constant depending only on \mathbf{G} . Such a bound follows (for instance) from the fact that this sum is a multiplicative character sum over the \mathbf{F}_ℓ -rational points of the geometrically connected affine algebraic variety $\mathbf{G}_{i,j}$ of dimension d .

Instead of looking for an elementary proof (which may well exist), we invoke the powerful ℓ -adic cohomological formalism (see, e.g. [67, 11.1.1] for an introduction, and compare with the proof of Proposition 8.8). Using the (rank 1) Lang–Kummer sheaf

$$\mathcal{K} = \mathcal{L}\left(\frac{g_{i,j}}{\ell}\right)$$

(over some p -adic field with $p \neq \ell$), we have by the Grothendieck–Lefschetz Trace Formula

$$\sum_{g \in \mathbf{G}_{i,j}(\mathbf{F}_\ell)} \left(\frac{g_{i,j}}{\ell} \right) = \sum_{g \in \mathbf{G}_{i,j}(\mathbf{F}_\ell)} \text{Tr}(\text{Fr}_{g,\ell} \mid \mathcal{K}) = \sum_{k=0}^{2d} (-1)^k \text{Tr}(\text{Fr} \mid H_c^k(\overline{\mathbf{G}}_{i,j}, \mathcal{K}))$$

where $\text{Fr}_{g,\ell}$ (respectively Fr) is the local (respectively global) geometric Frobenius for g seen as defined over \mathbf{F}_ℓ (resp. acting on the cohomology of the base-changed variety to an algebraic closure of \mathbf{F}_ℓ). By Deligne’s Riemann Hypothesis (see, e.g., [67, Theorem 11.37]), we have

$$\begin{aligned} \sum_{g \in \mathbf{G}_{i,j}(\mathbf{F}_\ell)} \left(\frac{g_{i,j}}{\ell} \right) &\ll \ell^d \dim H_c^{2d}(\overline{\mathbf{G}}_{i,j}, \mathcal{K}) + \ell^{d-1/2} \sum_{k < 2d} \dim H_c^k(\overline{\mathbf{G}}_{i,j}, \mathcal{K}) \\ &\ll \ell^d \dim H_c^{2d}(\overline{\mathbf{G}}_{i,j}, \mathcal{K}) + \ell^{d-1/2} \end{aligned}$$

for $\ell \geq 3$, by results of Bombieri or Adolphson–Sperber that show that the sum of dimensions of cohomology groups is bounded independently of ℓ (see, e.g., [67, Theorem 11.39]).

It therefore remains to prove that $H_c^{2d}(\overline{\mathbf{G}}_{i,j}, \mathcal{K}) = 0$. However, this space is isomorphic (as vector space) to the space of coinvariants of the geometric fundamental group of $\mathbf{G}_{i,j}$ acting on a one-dimensional space through the character which ‘is’ the Lang–Kummer sheaf \mathcal{K} . This means that either the coinvariant space is zero, and we are done, or otherwise the sheaf is geometrically trivial. The latter translates to the fact that the traces on \mathcal{K} of the local Frobenius Fr_{g,ℓ^v} of rational points $g \in \mathbf{G}_{i,j}(\mathbf{F}_{\ell^v})$ over all extensions fields $\mathbf{F}_{\ell^v}/\mathbf{F}_\ell$ depend only on v , i.e., the map

$$g \mapsto \left(\frac{N_{\mathbf{F}_{\ell^v}/\mathbf{F}_\ell} g_{i,j}}{\ell} \right)$$

on $\mathbf{G}_{i,j}(\mathbf{F}_{\ell^v})$ depends only on v . But this is clearly impossible for $SL(n)$ or $Sp(2g)$ with $n \geq 2$, $g \geq 1$ and $\ell \geq 3$, because if $\ell \geq 3$ we can explicitly write down matrices even in $\mathbf{G}(\mathbf{F}_\ell)$ both with $g_{i,j}$ a non-zero square and $g_{i,j}$ not a square.

(3) and (4): these are similar to (1). Namely, define now (again as in Chapter 8) a q -symplectic polynomial f in $\mathbf{F}_\ell[X]$ to be one of degree $2g$ such that³

$$f(0) = 1, \quad \text{and} \quad q^g T^{2g} f\left(\frac{1}{qT}\right) = f(T).$$

We can express such a q -symplectic polynomial uniquely in the form

$$\begin{aligned} f(T) = & 1 + a_1(f)T + \cdots + a_{g-1}(f)T^{g-1} + a_g(f)T^g \\ & + qa_{g-1}(g)T^{g+1} + \cdots + q^{g-1}a_1(f)T^{2g-1} + q^g T^{2g}, \end{aligned}$$

with $a_i(f) \in \mathbf{F}_\ell$, and this expression gives a bijection

$$f \mapsto (a_1(f), \dots, a_g(f))$$

between the set of q -symplectic polynomials and \mathbf{F}_ℓ^g .

Since the reversed characteristic polynomial $\det(1 - Tg)$ of a matrix $g \in CSp(2g, \mathbf{F}_\ell)$ is q -symplectic with $m(g) = q$, we need to bound

$$\frac{1}{|Sp(2g, \mathbf{F}_\ell)|} \sum_{f \in \Omega^{(\gamma)}} |\{g \in G \mid \det(1 - Tg) = f\}| \quad (\text{B.12})$$

where we have put (in Case (3) and (4), respectively)

$$\Omega^{(3)} = \{f \in \mathbf{F}_\ell[T] \mid f \text{ is } q\text{-symplectic and } f(1) \text{ is a square in } \mathbf{F}_\ell\},$$

$$\Omega^{(4)} = \{f \in \mathbf{F}_\ell[T] \mid f \text{ is } q\text{-symplectic and } q + 1 - a_1(f) \text{ is a square in } \mathbf{F}_\ell\}.$$

³ Unfortunately, this is not stated correctly in [80], although none of the results there are affected by this slip...

Now it is easy to check that we have

$$|\Omega^{(\gamma)}| = \frac{\ell^g + \ell^{g-1}}{2} \geq \frac{\ell^g}{2} \quad (\text{B.13})$$

for $\gamma = 3$ or 4 (recall ℓ is odd). Indeed, treating the case $\gamma = 3$ (the other is similar), we have

$$|\Omega^{(3)}| = |\{f \mid f(1) = 0\}| + \frac{1}{2} \sum_{f(1) \neq 0} \left(1 + \left(\frac{f(1)}{\ell}\right)\right).$$

The first term is ℓ^{g-1} since $f \mapsto f(1)$ is a non-zero linear functional on \mathbf{F}_ℓ^g . The first part of the second sum is $(\ell^g - \ell^{g-1})/2$, and the last is

$$\sum_{(a_2, \dots, a_{g-1})} \sum_{a_g \neq -\tilde{f}(1)} \left(\frac{a_g + \tilde{f}(1)}{\ell}\right),$$

where $\tilde{f}(1)$ is defined by $f(1) = a_g + \tilde{f}(1)$ (note that $\tilde{f}(1)$ depends only on the first variables (a_2, \dots, a_{g-1})). Because of the summation over the free variable a_g , this expression vanishes.

Now appealing to Lemma B.5, we obtain

$$\frac{1}{|Sp(2g, \mathbf{F}_\ell)|} |\{g \in G \mid \det(1 - Tg) = f\}| \geq \frac{1}{\ell^g} \left(\frac{\ell}{\ell + 1}\right)^{2g^2 + g + 1} \quad (\text{B.14})$$

for all q -symplectic polynomials f , and hence the stated bound follows by combining (B.12), (B.13), and (B.14). (Note that the two definitions of symplectic polynomials correspond via the relation between the characteristic polynomial and the reversed characteristic polynomial, so counting one type or the other is equivalent.)

(5) and (6): this is again similar to (3) and (4), where we now deal with

$$\frac{1}{|Sp(2g, \mathbf{F}_\ell)|} \sum_{f \in \Omega^{(\gamma)}} |\{g \in G \mid \det(1 - Tg) = f\}|$$

with $\gamma = 5$ or 6 and

$$\Omega^{(5)} = \{f \in \mathbf{F}_\ell[T] \mid f \text{ is } q\text{-symplectic and } f(1) = 0\},$$

$$\Omega^{(6)} = \{f \in \mathbf{F}_\ell[T] \mid f \text{ is } q\text{-symplectic and } q + 1 = a_1(f)\}.$$

We have in both cases $|\Omega^{(\gamma)}| = \ell^{g-1}$ since the condition is a linear one on the coefficients. By Lemma B.5 (and the same remark as before), we also have

$$|\{g \in G \mid \det(1 - Tg) = f\}| \leq \frac{1}{\ell^g} \left(\frac{\ell}{\ell - 1}\right)^{2g^2 + g + 1}$$

for all f , and therefore

$$\frac{1}{|Sp(2g, \mathbf{F}_\ell)|} \sum_{f \in \Omega'} |\{g \in G \mid \det(1 - Tg) = f\}| \leq \frac{1}{\ell} \left(\frac{\ell}{\ell - 1} \right)^{2g^2 + g + 1}.$$

Moreover, the quantity to estimate is also at most 1 for trivial reasons, and hence the stated bound follows.

(7) This final density result is due to Dunfield–Thurston [34, 8.2, 8.3], and we include a sketch of the proof for completeness. Since all alternating forms are equivalent, and the symplectic group acts transitively on the set of all Lagrangians, it suffices to work with the ‘model’ mentioned in the Section on notation: $J = \mathbf{F}_\ell^g \subset V = J \oplus J'$ where the symplectic form is

$$\langle (v_1, \ell_1), (v_2, \ell_2) \rangle = \ell_1(v_2) - \ell_2(v_1).$$

Moreover, again because the action on Lagrangians is transitive, the desired density is equal to $|\mathcal{L}^*|/|\mathcal{L}|$, where \mathcal{L} is the set of Lagrangians in V , and $\mathcal{L}^* \subset \mathcal{L}$ is the set of those which are transverse to J .

First, one computes $|\mathcal{L}|$. Using the transitive action of $Sp(V)$, it suffices to know the order of the stabilizer H of J in $Sp(V)$. This is determined by looking at the natural homomorphism

$$\varphi : H \rightarrow GL(J)$$

which is surjective (because of the section $x \mapsto x \oplus (x^*)^{-1}$) and has kernel in bijection with the set \mathcal{S} of ‘symmetric endomorphisms’ $J' \rightarrow J$, i.e., those A such that

$$\langle A(\ell_1), \ell_2 \rangle = \langle A(\ell_2), \ell_1 \rangle;$$

indeed, this bijection is given by

$$\begin{cases} \text{Ker}(\varphi) & \rightarrow & GL(J') \\ x & \mapsto & A, \end{cases}$$

where x restricted to J' is the direct sum of $A : J' \rightarrow J$ and $A' : J' \rightarrow J'$. Therefore, we have

$$|\mathcal{L}| = \frac{|Sp(2g, \mathbf{F}_\ell)|}{|GL(J)| \cdot |\mathcal{S}|}.$$

It remains to compute $|\mathcal{L}^*|$. For this, notice that each transverse Lagrangian is the graph of a linear map $A : J' \rightarrow J$, and conversely that the graph of a linear map A defines a (unique) Lagrangian, if and only if it is symmetric in the sense above. Hence, again, the number of A is $|\mathcal{S}|$, so

$$\frac{|\mathcal{L}^*|}{|\mathcal{L}|} = \frac{|\mathcal{S}|^2}{|GL(g, \mathbf{F}_\ell)| |Sp(2g, \mathbf{F}_\ell)|}.$$

Since $|\mathcal{S}| = \ell^{g(g+1)/2}$ (taking a basis of J and the dual basis of J' , \mathcal{S} is in bijection correspondence with symmetric $g \times g$ matrices), applying (0.1), (0.2) concludes the proof. \square

B.3 Other techniques

Finding the local densities of sieving sets Ω_ℓ , when those are defined ‘over \mathbf{F}_ℓ ’ as subsets of \mathbf{F}_ℓ^r , or of the rational points of some more general algebraic variety, may be quite a challenge.

In addition to the results of the previous sections, which are quite versatile in their way, we want to point out another technique that can be useful.

Suppose that $Y_\ell \subset \mathbf{F}_\ell^r$ for definiteness. Then quite often, even if Ω_ℓ is not the set of points of an algebraic variety over \mathbf{F}_ℓ (i.e., defined as the set of solutions of some polynomial equations), it may have the form of a *definable set*, in the sense of logic, in the first-order language of rings. Without writing down the full formal definition (see, e.g., [83, Section 2]), this essentially means that Ω_ℓ may be defined using not only conditions of the type $f(x_1, \dots, x_r) = 0$, where f is some polynomial, but also by negations $f(x_1, \dots, x_r) \neq 0$, and by combinations of such elementary terms using the logic connectors ‘and’, ‘or’, ‘implies’, and quantified expressions over variables *ranging over* \mathbf{F}_ℓ . So for instance, the set Ω_ℓ of irreducible monic polynomials of degree 2 can be defined as the set of $(a, b) \in \mathbf{F}_\ell^2$ for which the following logical formula is true:

$$(\forall x, x^2 + ax + b \neq 0).$$

More generally, the sets of polynomials and matrices in Lemma B.1 and Proposition B.4 are all definable subsets over finite fields in this sense (using the obvious variables, either coefficients of polynomials or of matrices).

Now it turns out that, although the definition may look rather more complicated than that of an algebraic variety, the cardinality of such definable sets is quite well-behaved. Indeed, we have the following remarkable result of Chatzidakis, van den Dries, and Macintyre [21]:

Theorem B.6 *Let $\varphi(x, y)$ be a formula in the language of rings with variables (x_1, \dots, x_n) and parameters (y_1, \dots, y_m) . There exist a prime power q_0 , a constant $C \geq 0$ depending only on φ and q_0 , and a finite family (d_i, δ_i) of pairs where d_i is an integer with $0 \leq d_i \leq n$, and $\delta_i > 0$ is a rational number, such that for any prime power $q \geq q_0$, any $y = (y_1, \dots, y_m) \in \mathbf{F}_q^m$, either the set*

$$\varphi(\mathbf{F}_q, y) = \{x \in \mathbf{F}_q^n \mid \varphi(x, y) \text{ is true}\}$$

is empty, or it satisfies

$$|\varphi(\mathbf{F}_q, y) - \delta_i q^{d_i}| \leq C q^{d_i-1/2}$$

for some i .

Intuitively, δ_i is a ‘density’ and d_i a ‘dimension’; the fact that those may vary with q is not surprising since this happens already for algebraic varieties (e.g., the variety defined by $X^2 + 1 = 0$). Note the additional parameters which indicate the great uniformity of such estimates.

See also [83] for a potentially useful study of exponential sums over such definable sets.

Appendix C

Representation theory

This Appendix quickly surveys some aspects of representation theory that we use in this book, highlighting the aspects which are most relevant. We refer to [115] for a more complete treatment of the case of finite groups.

C.1 Definitions

A *linear representation* of a group G , defined over a field K , is a group homomorphism

$$\rho : G \rightarrow GL(V)$$

where V is a K -vector space. In other words, ρ ‘is’ an action of G on V by linear transformations, and we write $\rho(g)v$ or simply $g \cdot v$ for this action. If ρ is injective, the representation is called *faithful*. We also denote $V = V_\rho$ when we want the notation to be more specific. If $K = \mathbf{C}$, which is assumed unless otherwise specified, the representation is *unitary* if there exists an inner product on V (making it a Hilbert space), so that the operators $\rho(g)$ are in the unitary group $U(V)$ for all $g \in G$, or equivalently, if G acts on V by linear isometries for some inner product. Also, when G and $GL(V)$ carry a topology (compatible with the group structure), then the map ρ is assumed to be continuous.

If V is finite-dimensional, the *degree* or *dimension* of the representation is the dimension of V , denoted by $\dim \rho$. A simple but important example of representation is the map sending all $g \in G$ to $1 \in K^\times = GL(1, K)$; this is called the *trivial representation*.

Part of the importance of representations in general stems from their malleability, arising from the flexible formalism of linear algebra. In particular, one can define a *morphism* $\rho \rightarrow \tau$ between representations ρ and τ to be a K -linear map

$$V_\rho \xrightarrow{\varphi} V_\tau$$

such that for any $g \in G$, the square diagram

$$\begin{array}{ccc} V_\rho & \xrightarrow{\varphi} & V_\pi \\ \rho(g) \downarrow & & \downarrow \tau(g) \\ V_\rho & \xrightarrow{\varphi} & V_\pi \end{array} \quad (\text{C.1})$$

commutes (i.e., for all $g \in G$ and $v \in V_\rho$, we have $\varphi(\rho(g)v) = \tau(g)\varphi(v)$, or in shorthand, $\varphi(g \cdot v) = g \cdot \varphi(v)$). Morphisms are *injective*, *surjective* or are *isomorphisms* if φ has the corresponding property as a linear map; of course, if φ is an isomorphism, its inverse is also a morphism $\tau \rightarrow \rho$. This is then denoted $\rho \simeq \tau$.

Moreover, linear algebra operations provide means of constructing new representations from existing ones: the *direct sum* $\rho \oplus \tau$ is defined as the representation mapping g to $\rho(g) \oplus \tau(g)$ acting on $V_\rho \oplus V_\tau$ componentwise; the *tensor product* $\rho \otimes \tau$ is defined as the representation mapping g to $\rho(g) \otimes \tau(g)$ acting on $V_\rho \otimes V_\tau$. When defined, the dimensions are given by

$$\dim(\rho \oplus \tau) = \dim \rho + \dim \tau, \quad \dim(\rho \otimes \tau) = (\dim \rho)(\dim \tau).$$

More generally, the tensor product can be used to define representations of a direct product $G_1 \times G_2$: if ρ is a representation of G_1 and τ is one of G_2 , then

$$(g_1, g_2) \mapsto \rho(g_1) \otimes \tau(g_2)$$

is a representation of $G_1 \times G_2$ on the space $V_\rho \otimes V_\tau$, denoted $\rho \boxtimes \tau$, and called the *external tensor product* of ρ and τ . Note that if $G = G_1 = G_2$, the composite

$$\begin{cases} G & \longrightarrow & G \times G & \xrightarrow{\rho \boxtimes \tau} & V_\rho \otimes V_\tau \\ g & \mapsto & (g, g) & & \end{cases}$$

yields the usual tensor product $\rho \otimes \tau$.

Also important is the *contragredient* $\tilde{\rho}$ or $\bar{\rho}$ of a representation ρ , which acts on the dual space $V_{\bar{\rho}} = V'_\rho = \text{Hom}_K(V_\rho, K)$ by

$$(\bar{\rho}(g)\ell)(v) = \ell(\rho(g^{-1})v), \quad \text{for all linear forms } \ell \in V'_\rho \text{ and } v \in V_\rho,$$

or equivalently with the duality bracket $\langle \ell, v \rangle = \ell(v)$, we have the shorthand

$$\langle g \cdot \ell, v \rangle = \langle \ell, g^{-1} \cdot v \rangle.$$

As an example of isomorphism, it is easy to see that if ρ is finite-dimensional, we have

$$\rho \otimes \bar{\rho} \simeq \text{End}(V_\rho)$$

where G acts on the space of K -linear endomorphisms of V_ρ by

$$g \cdot A = \rho(g)^{-1} A \rho(g), \quad \text{for all } A \in \text{End}(V_\rho);$$

this isomorphism is given by the map $v \otimes \ell \mapsto (w \mapsto \ell(w)v)$ (which gives the standard K -linear isomorphism between $V \otimes V'$ and $\text{End}(V)$ for a finite-dimensional K -vector space).

The last important operations concerning representations we will use are restriction and induction. The first is quite clear: if H is a subgroup of G , then any representation ρ of G restricts to one of H , which is denoted by $\text{Res}_H^G(\rho)$. On the other hand, defining induction is not so obvious; assume that H is of finite index in G , and let ρ be a representation of H . Then defining

$$W = \{f : G \rightarrow V_\rho \mid f(hx) = \rho(h)(f(x)) \text{ for all } h \in H, x \in X\},$$

we obtain a K -vector space, on which G acts by $g \cdot f(x) = f(xg)$ for $f \in W$ and $x \in G$. The corresponding representation is called the *representation induced to G by ρ* , and is denoted by $\text{Ind}_H^G(\rho)$. The specific construction of W is not important, and in fact the following important relation (which is one form of *Frobenius reciprocity*) is often the only information required about induced representations: for any representation ρ of G and τ of H , we have

$$\text{Hom}_H(\text{Res}_H^G(\rho), \tau) \simeq \text{Hom}_G(\rho, \text{Ind}_H^G(\tau)), \quad (\text{C.2})$$

where $\text{Hom}_G(\cdot, \cdot)$ is the space of morphisms *as representations* between two representations of a group G . Using the description of the induced representation $\text{Ind}_H^G(\tau)$ on the space W above, this map is obtained as follows: given an H -homomorphism $\varphi : V_\rho \rightarrow V_\tau$, its image is $\tilde{\varphi} : V_\rho \rightarrow W$ such that $\tilde{\varphi}(v)$ is the function $g \mapsto \varphi(gv)$, which lies in W by the assumption that φ commutes with the action of H .

In this book, the representations which occur are finite-dimensional and (except partly in Chapter 8) have the further property that ρ factors through a finite quotient of G , i.e., $\text{Ker } \rho$ is of finite index in G . This condition implies in particular that those representations which are defined over \mathbf{C} are always unitary: if $\langle \cdot, \cdot \rangle$ is an arbitrary inner product on V_ρ , then we can define another G -invariant inner product by putting

$$\langle v, w \rangle_\rho = \frac{1}{|\text{Im}(\rho)|} \sum_{T \in \text{Im}(\rho)} \langle T(v), T(w) \rangle$$

since this is a finite sum, and the averaging has the obvious effect that

$$\langle \rho(g)v, \rho(g)w \rangle_\rho = \langle v, w \rangle_\rho$$

as desired. A practical consequence is that all eigenvalues of $\rho(g)$, $g \in G$, are roots of unity. Observe also that if ρ is isomorphic to a direct sum $\rho_1 \oplus \rho_2$, then we can always find an invariant inner product so that the direct sum $V_{\rho_1} \oplus V_{\rho_2}$ is orthogonal.

C.2 Harmonic analysis

Representation theory is a vast subject and representations serve many purposes. One of our primary interests in this book is that representations provide a tool to efficiently analyze functions defined on G . Indeed, suppose G is finite; then given a finite-dimensional representation ρ (over \mathbf{C}), we can define a function

$$\chi_\rho \begin{cases} G \rightarrow \mathbf{C} \\ x \mapsto \text{Tr } \rho(x), \end{cases}$$

which is called the *character* of ρ , and if ρ is unitary (which we can always assume), we can define many functions (called *matrix coefficients*) by choosing vectors $v, w \in V_\rho$ and defining

$$\varphi_{v,w} \begin{cases} G \rightarrow \mathbf{C} \\ x \mapsto \langle \rho(x)v, w \rangle \end{cases}$$

using the G -invariant inner product on V_ρ . Note the simple but useful bounds

$$|\chi_\rho(x)| \leq \chi_\rho(1) = \dim \rho, \quad |\varphi_{v,w}(x)| \leq \|v\| \|w\| \dim \rho, \quad \text{for all } x \in G$$

since the eigenvalues of $\rho(x)$ are roots of unity.

The first main point is that those functions can be used to generate the space of functions on G , as \mathbf{C} -vector space. More precisely, observe that this space, which we denote $L^2(G)$, is itself a finite-dimensional Hilbert space by means of the inner product

$$\langle f, g \rangle = \frac{1}{|G|} \sum_{x \in G} f(x) \overline{g(x)}, \quad (\text{C.3})$$

so it is also desirable to have an orthonormal basis of $L^2(G)$.

If we look at characters first, it is clear that they can not span all of $L^2(G)$ because, being defined as traces, they are functions invariant under conjugation: we have $\chi_\rho(yxy^{-1}) = \chi_\rho(x)$ for all $x, y \in G$. Denote by $L^2(G^\#)$ the subspace of *class functions*, which are those satisfying this relation, with the induced inner product. Now observe that if we want to use characters of representations to generate minimally the space of class functions, there are some obvious redundancies: on the one hand, it is also clear from the diagram (C.1) that $\chi_\rho = \chi_\tau$ if $\rho \simeq \tau$, so we need only keep one representative of each isomorphism class of representations; on the other hand, we have equally obviously $\chi_{\rho \oplus \tau} = \chi_\rho + \chi_\tau$, so that whenever a representation is (isomorphic to) the direct sum of at least two representations, we need only keep the characters of those components.

Clearly, this means that representations which *cannot* be so decomposed play a crucial role. They are called *irreducible representations*, and are characterized as those representations ρ which have no non-trivial G -invariant subspace.

This is not a trivial fact (and is false for infinite groups in general), since it entails showing that any such subspace W has a G -invariant complementary subspace, so that ρ is the direct sum of the representation of G induced on W and another representation on this complement (in other words, any representation is *completely reducible*, or *semisimple*). In the case considered, this is easily seen from the unitarity of ρ : the orthogonal complement W^\perp of W with respect to a G -invariant inner product will be itself invariant under G , and gives the required decomposition into a direct sum.

Now the first important result we have about representations is the following:

Proposition C.1 *The distinct characters χ_ρ of the irreducible representations of a finite group G form an orthonormal basis of the space $L^2(G^\sharp)$ of class functions on G .*

These characters are naturally called *irreducible characters*.

Note in particular that this proposition means that if two irreducible representations have the same character (as function on G), they must be isomorphic, which is by no means obvious. Moreover, the number of irreducible representations of G , up to isomorphism, is equal to the dimension of $L^2(G^\sharp)$, and therefore is the same as the number of conjugacy classes in G .

If $G = G_1 \times G_2$ is a direct product then it is not hard to deduce from this proposition that all irreducible representations of G are of the form $\rho \boxtimes \tau$ for some irreducible representations ρ of G_1 and τ of G_2 , and that the elements of the basis of $L^2(G^\sharp)$ are given by

$$(g_1, g_2) \mapsto \chi_\rho(g_1)\chi_\tau(g_2)$$

where (χ_ρ) is the basis for G_1 and (χ_τ) that for G_2 .

Although this proposition encapsulates one important use of representations (and the one most relevant for our general description of conjugacy sieves on groups), it is really a reflection of a more algebraic phenomenon. Namely, the space $L^2(G)$ itself is a representation, called the *regular representation*¹ of G by the rule

$$\rho(g)f(x) = f(xg)$$

for $f \in L^2(G)$, $g \in G$ (note that on the other hand, $L^2(G^\sharp)$ is not in general a representation of G in a natural way). This representation is unitary with respect to the inner product (C.3). Then Proposition C.1 is related to the fact that $L^2(G)$ is isomorphic to the orthogonal direct sum

¹ Not to be confused with the regular characters of finite groups of Lie type that occur in Chapter 5.

$$L^2(G) = \bigoplus_{\rho} (\dim \rho) \rho, \quad (\text{C.4})$$

where ρ runs over all isomorphism classes of irreducible representations of G , and where $n\rho$, for $n \geq 1$ an integer, is shorthand for an orthogonal direct sum of n representations, each of which is isomorphic to ρ . This means in particular that *every* irreducible representation occurs as a sub-representation of $L^2(G)$, and that we have the relation

$$\dim G = \sum_{\rho} (\dim \rho)^2.$$

Now the existence of this decomposition is not a special property of $L^2(G)$: any finite-dimensional representation τ of G can be decomposed as an orthogonal direct sum

$$\tau \simeq \bigoplus_{\rho} m_{\rho}(\tau) \rho,$$

where ρ runs again over irreducible representations of G up to isomorphism, for some integers $m_{\rho}(\tau) \geq 0$, which are called the *multiplicities* in τ of the representations ρ . Moreover, those multiplicities are uniquely determined, so that the decomposition is unique in an obvious sense. In fact, it can be determined concretely by means of the formulas

$$m_{\rho}(\tau) = \dim \operatorname{Hom}_G(\tau, \rho) = \langle \chi_{\tau}, \chi_{\rho} \rangle = \frac{1}{|G|} \sum_{x \in G} \operatorname{Tr}(\tau(x)) \overline{\operatorname{Tr}(\rho(x))},$$

for the multiplicities. This again explains the importance of the characters of irreducible representations since, knowing them and the character of any representation, we can decompose the latter into a sum of irreducibles.

It follows also that an arbitrary representation τ is determined, up to (non-unique) isomorphism, by its character, since the multiplicities are. In addition, for any representations τ_1 and τ_2 , we have

$$\dim \operatorname{Hom}_G(\tau_1, \tau_2) = \langle \chi_{\tau_1}, \chi_{\tau_2} \rangle$$

(by linearity from the multiplicity formula, for example).

It is often very useful to identify an isomorphism class of representations with its character. In fact, it is also extremely convenient to consider class functions on G which are linear combinations with integral, but not necessarily positive, coefficients, e.g., $f = \chi_1 - \chi_2$ with χ_i an irreducible character. Such *generalized characters* are fundamental to the Deligne–Lusztig theory of representations of finite matrix groups, which is used in Chapter 5. They form a free abelian group of finite type, and the irreducible characters form a basis of it.

With this identification, one then writes the Frobenius reciprocity formula in the form

$$\langle \text{Res}_H^G(\rho), \tau \rangle = \langle \rho, \text{Ind}_H^G(\tau) \rangle. \quad (\text{C.5})$$

As an example of application of this formula, note that from it one recovers immediately the decomposition (C.4) of the regular representation: indeed, the definition itself shows that $L^2(G) \simeq \text{Ind}_1^G(1)$, i.e., the regular representation is induced from the trivial representation of the trivial subgroup, and hence for any irreducible representation ρ of G , we have

$$\langle L^2(G), \rho \rangle = \langle \text{Ind}_1^G(1), \rho \rangle = \langle 1, \text{Res}_1^G(\rho) \rangle = \dim \rho.$$

Since characters of direct sums are sums of characters, and since we have

$$\chi_{\rho \otimes \tau} = \chi_\rho \chi_\tau, \quad \chi_{\bar{\rho}} = \overline{\chi_\rho},$$

the formula for the multiplicity easily gives relations such as

$$\langle \rho_1 \otimes \rho_2, \rho_3 \rangle = \langle \rho_1, \overline{\rho_2} \otimes \rho_3 \rangle, \quad \text{etc.}$$

An important special case is the multiplicity $m_1(\tau)$ of the trivial representation in a representation τ : this is none other than the dimension of the space V_τ^G of vectors which are invariant under G . In particular, the multiplicity formula states that τ has no (non-zero) invariant vector if and only if the character of τ is orthogonal to $\chi_1 = 1$, i.e., if the average value of χ_τ on G is zero.

Coming back to the harmonic analysis, Proposition C.1 gives a good natural basis of $L^2(G^\sharp)$. There is no such intrinsic basis of the whole of $L^2(G)$, in general, but one can still be fairly explicit:

Proposition C.2 *For each irreducible representation ρ of G , let*

$$(e_{\rho,1}, \dots, e_{\rho,\dim \rho})$$

be an arbitrary fixed orthonormal basis of V_ρ with respect to a G -invariant inner product. Then the family of matrix coefficients

$$g \mapsto \sqrt{(\dim \rho)} \langle \rho(g) e_{\rho,i}, e_{\rho,j} \rangle$$

as ρ runs over irreducible representations of G and $1 \leq i, j \leq \dim \rho$, is an orthonormal basis of $L^2(G)$.

C.3 One-dimensional representations

Among the representations, those of dimension 1 are somewhat special and easier to deal with. Indeed, since $GL(1, K)$ is abelian, we have the usual bijection

$$\mathrm{Hom}(G, GL(1, K)) \simeq \mathrm{Hom}(G/[G, G], GL(1, K))$$

which shows that representations of G of dimension 1 are in one-to-one correspondence with those of the abelianization $G^{ab} = G/[G, G]$.

So the study of representations of dimension 1 reduces to the case of an abelian group G . Then the situation simplifies further for a number of reasons: first, all irreducible representations of a finite abelian group are indeed of dimension 1, and they really are ‘the same’ as their character; second, irreducible representations of dimension 1 themselves form a group (called the *character group* of G) in a natural way, namely

$$(\chi_1 \chi_2)(g) = \chi_1(g) \chi_2(g);$$

finally, there is no difference between class functions and arbitrary functions, so that characters can be used to easily expand any function on G in terms of the distinguished basis of characters, which is the basis of functions $\chi : G \rightarrow \mathbf{C}$ satisfying $\chi(xy) = \chi(x)\chi(y)$ for all $x, y \in G$.

For the reasons above, it is customary to simply speak of a *character* of an abelian group, meaning an irreducible representation seen as a function $G \rightarrow \mathbf{C}^\times$.

C.4 The character tables of $GL(2, \mathbf{F}_q)$ and $SL(2, \mathbf{F}_q)$

According to Proposition C.1 and the general decomposition formula, much of the representation theory of a finite group G is accessible in principle if the characters of G , as functions on G , or equivalently, on the set of conjugacy classes of G , are explicitly known. This data is called the *character table* of G , because of the way it is naturally presented as a table listing character values at each conjugacy class.

The character tables of the ‘simplest’ finite groups of Lie type, $GL(2, \mathbf{F}_q)$ and $SL(2, \mathbf{F}_q)$, q a power of a prime, are good illustrations both of the general theory and of the theory of Deligne–Lusztig characters which we use in Chapter 5. Those particular tables (due to Frobenius) are found in almost all textbooks (e.g. in [31, 15.9], or in [44, p. 70]), but we include them for the reader’s convenience. We label the representations according to the Deligne–Lusztig terminology, which is briefly explained in Chapter 5; textbooks which do not cover this theory will have different notation.

First of all we list the conjugacy classes in Table C.1. For this we fix an element $\varepsilon \in \mathbf{F}_q^\times$ which is not a square, so that $\mathbf{F}_q(\sqrt{\varepsilon}) = \mathbf{F}_{q^2}$. The third column of this table lists which further conjugacies hold in a given line.

Table C.1 Conjugacy classes of $GL(2, \mathbf{F}_q)$

Form	Condition	Equivalence	Number	Cardinality
$\begin{pmatrix} x & 0 \\ 0 & x \end{pmatrix}$	$x \in \mathbf{F}_q^*$	None	$q - 1$	1
$\begin{pmatrix} x & 0 \\ 0 & y \end{pmatrix}$	$x \neq y, xy \neq 0$	$(x, y) \sim (y, x)$	$\frac{1}{2}(q - 1)(q - 2)$	$q(q + 1)$
$\begin{pmatrix} a & b \\ \varepsilon b & a \end{pmatrix}$	$b \neq 0$	$(a, b) \sim (a, -b)$	$\frac{1}{2}q(q - 1)$	$q(q - 1)$
$\begin{pmatrix} a & 1 \\ 0 & a \end{pmatrix}$	$a \neq 0$	None	$q - 1$	$q^2 - 1$

Note that there are $q^2 - 1$ conjugacy classes. Now we list the character table, starting with a list of the different types of representations, with their number and dimensions. In Table C.2, χ, χ_1, χ_2 are characters of \mathbf{F}_q^\times and ψ is a character of $\mathbf{F}_{q^2}^\times$. The third column indicates the isomorphisms to be taken into account. Moreover, in terms of the element ε previously defined such that $\sqrt{\varepsilon}$ generates \mathbf{F}_{q^2} , we write

$$\alpha = a + \sqrt{\varepsilon}b, \quad \bar{\alpha} = a - \sqrt{\varepsilon}b, \quad N\alpha = a^2 - \varepsilon b^2.$$

In the notation of Deligne–Lusztig characters, the two types of maximal rational tori are represented by

$$\mathbf{T} = \left\{ \begin{pmatrix} x & 0 \\ 0 & y \end{pmatrix} \right\}, \quad \mathbf{T}_s = \left\{ \begin{pmatrix} a & b \\ \varepsilon b & a \end{pmatrix} \right\},$$

where $x, y \in \mathbf{F}_q^\times$ and $(a, b) \in \mathbf{F}_q^2 - \{(0, 0)\}$; the first one is *split* and the second one is not, and is obtained from the split torus by twisting with the non-trivial element

$$s = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

of the Weyl group of $GL(2)$.

The character table properly speaking is in Table C.3.

In the usual terminology, the irreducible representations $R_{\mathbf{T}}^G(\chi_1, \chi_2)$ are called the *principal series*, and the $-R_{\mathbf{T}_s}^G(\psi)$ are called the *discrete series*.

Note that, with ρ running over irreducible representations, we have

$$\sum_{\rho} \dim(\rho) = q^3 - q^2, \quad \max_{\rho} \dim(\rho) = \begin{cases} q + 1, & \text{if } q > 2 \\ 2, & \text{if } q = 2. \end{cases}$$

Table C.2 Irreducible representations of $GL(2, \mathbf{F}_q)$

Type	Condition	Isomorphisms	Number	Dimension
$\chi \circ \det$	None	None	$q - 1$	1
$\text{St} \circ \det$	None	None	$q - 1$	q
$R_T^G(\chi_1, \chi_2)$	$\chi_1 \neq \chi_2$	$(\chi_1, \chi_2) \sim (\chi_2, \chi_1)$	$\frac{1}{2}(q-1)(q-2)$	$q+1$
$-R_{T_s}^G(\psi)$	$\psi \neq \psi^q$	$\psi \sim \psi^q$	$\frac{1}{2}q(q-1)$	$q-1$

Table C.3 Character table of $GL(2, \mathbf{F}_q)$

	$\begin{pmatrix} x & 0 \\ 0 & x \end{pmatrix}$	$\begin{pmatrix} x & 0 \\ 0 & y \end{pmatrix}$	$\begin{pmatrix} a & b \\ \varepsilon b & a \end{pmatrix}$	$\begin{pmatrix} a & 1 \\ 0 & a \end{pmatrix}$
$\chi \circ \det$	$\chi(x^2)$	$\chi(xy)$	$\chi(N\alpha)$	$\chi(a^2)$
$R_T^G(\chi_1, \chi_2)$	$(q+1)\chi_1(x)\chi_2(x)$	$\chi_1(x)\chi_2(y) + \chi_2(x)\chi_1(y)$	0	$\chi_1(a)\chi_2(a)$
$-R_{T_s}^G(\psi)$	$(q-1)\psi(x)$	0	$-\psi(\alpha) - \psi(\bar{\alpha})$	$-\psi(a)$
$\text{St} \circ \chi$	$q\chi(x^2)$	$\chi(xy)$	$-\chi(N\alpha)$	0

Table C.4 Conjugacy classes of $SL(2, \mathbf{F}_q)$, q odd

Form	Condition	Equivalence	Number	Cardinality
$\begin{pmatrix} x & 0 \\ 0 & x \end{pmatrix}$	$x = \pm 1$	None	2	1
$\begin{pmatrix} x & 0 \\ 0 & x^{-1} \end{pmatrix}$	$x \neq 0, \pm 1$	$x \sim x^{-1}$	$\frac{1}{2}(q-3)$	$q(q+1)$
$\begin{pmatrix} a & b \\ \varepsilon b & a \end{pmatrix}$	$N\alpha = 1, b \neq 0$	$(a, b) \sim (a, -b)$	$\frac{1}{2}(q-1)$	$q(q-1)$
$\begin{pmatrix} a & 1 \\ 0 & a \end{pmatrix}$	$a = \pm 1$	None	2	$\frac{1}{2}(q^2-1)$
$\begin{pmatrix} a & \varepsilon \\ 0 & a \end{pmatrix}$	$a = \pm 1$	None	2	$\frac{1}{2}(q^2-1)$

Now we consider $SL(2, \mathbf{F}_q)$ where q is odd. The list of conjugacy classes is in Table C.4.

The list of representations is in Table C.5, with the number of each type and their dimensions; again χ is a character of \mathbf{F}_q^\times and ψ is a character of \mathbf{F}_q^\times . The

Table C.5 Irreducible representations of $SL(2, \mathbf{F}_q)$, q odd

Type	Condition	Isomorphisms	Number	Dimension
1	None	None	1	1
St	None	None	1	q
$R_{\mathbf{T}}^G(\chi, 1)$	$\chi \neq 1$	$\chi \sim \chi^{-1}$	$\frac{1}{2}(q-3)$	$q+1$
$-R_{\mathbf{T}_s}^G(\psi)$	$\psi^2 \neq 1, \psi^q \neq \psi$	$\psi \sim \psi^q, \psi \sim \psi^{-1}$	$\frac{1}{2}(q-1)$	$q-1$
$R^\pm(\chi_2)$	None	None	2	$\frac{1}{2}(q+1)$
$-R^\pm(\psi_2)$	None	None	2	$\frac{1}{2}(q-1)$

Table C.6 Character table of $SL(2, \mathbf{F}_q)$, q odd

	$\begin{pmatrix} x & 0 \\ 0 & x \end{pmatrix}$	$\begin{pmatrix} x & 0 \\ 0 & x^{-1} \end{pmatrix}$	$\begin{pmatrix} a & b \\ \varepsilon b & a \end{pmatrix}$	$\begin{pmatrix} a & 1 \\ 0 & a \end{pmatrix}$	$\begin{pmatrix} a & \varepsilon \\ 0 & a \end{pmatrix}$
1	1	1	1	1	1
St	q	1	-1	0	0
$R_{\mathbf{T}}^G(\chi)$	$(q+1)\chi(x)$	$\chi(x) + \bar{\chi}(x)$	0	$\chi(a)$	$\chi(a)$
$-R_{\mathbf{T}_s}^G(\psi)$	$(q-1)\psi(x)$	0	$-\psi(\alpha) - \psi(\bar{\alpha})$	$-\psi(a)$	$-\psi(a)$
$R^\pm(\chi_2)$	$\frac{1}{2}(q+1)\chi_2(x)$	$\chi_2(x)$	0	$\frac{1}{2}(\chi_2(a) \pm \omega)$	$\frac{1}{2}(\chi_2(a) \mp \omega)$
$-R^\pm(\psi_2)$	$\frac{1}{2}(q-1)\psi_2(x)$	0	$-\psi_2(\alpha)$	$-\frac{a}{2}(\chi_2(a) \mp \omega)$	$-\frac{a}{2}(\chi_2(a) \pm \omega)$

tori \mathbf{T} and \mathbf{T}_s are obtained as the intersection of the ones for $GL(2, \mathbf{F}_q)$ with $SL(2, \mathbf{F}_q)$; the twisting element s is now

$$\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}.$$

Those representations arise for the most part as restrictions of irreducible representations of $GL(2, \mathbf{F}_q)$, the only exceptions being the four ‘exceptional’ representations denoted $R^\pm(\chi_2)$ and $-R^\pm(\psi_2)$, which are the two irreducible components of the restriction to $SL(2, \mathbf{F}_q)$ of $R_{\mathbf{T}}^G(\chi_2, 1)$ and $-R_{\mathbf{T}_s}^G(\psi_2)$ (respectively), where χ_2 (respectively ψ_2) is the non-trivial character of order 2 of \mathbf{F}_q^\times (respectively $\mathbf{F}_{q^2}^\times$):

$$R_{\mathbf{T}}^G(\chi_2, 1) = R^+(\chi_2) \oplus R^-(\chi_2), \quad -R_{\mathbf{T}_s}^G(\psi_2) = (-R^+(\psi_2)) \oplus (-R^-(\psi_2)).$$

The precise character table, where only the character values for those last four representations are not obvious from Table C.3, is in Table C.6. They involve the following further notation: $\chi_2(-1) \in \{\pm 1\}$ is the value of the quadratic character at -1 and $\omega \in \mathbf{C}$ is such that $\omega^2 = \chi_2(-1)q$.

One can write down explicitly and completely the character tables for some other finite groups of Lie type with small rank, before they become unwieldy; see for instance the character tables of $GL(3, \mathbf{F}_\ell)$ and $GL(4, \mathbf{F}_\ell)$ in [126].

Appendix D

Property (T) and Property (τ)

This Appendix is, first, a review of the definition of Property (T) and Property (τ) , together with some simple examples and properties. The basic reference we use here is [58]; see also [91], [93], and [9], which also contains a survey of many applications of Property (T) . In Section D.4, we give most of the proof of Property (T) for $SL(n, \mathbf{Z})$ for $n \geq 3$ due to Shalom [118], adapted to the simpler case of Property (τ) for finite-dimensional representations (this restriction avoids the use of the spectral theorem for infinite-dimensional representations of abelian groups, and it corresponds to the applications involving the large sieve).

D.1 Property (T)

Let G be a group. The set of all (continuous) unitary representations of G is usually badly understood and unwieldy, and lacking in ‘easy’ structural properties. Infinite-dimensional representations, in particular, can not be studied using their characters, since self-adjoint operators in infinite-dimensional Hilbert spaces do not usually have a well-defined trace. However, Kazhdan realized in the 1960s that certain useful properties existed that could be proved *independently* of a precise knowledge of the set of all representations.

We assume that G is locally compact (for some topology for which the group law and inverse are continuous). Then G has *Property (T)* or *is a Kazhdan group* if there exists a compact subset $K \subset G$ and $\varepsilon > 0$ such that for any continuous unitary representation ρ of G on a Hilbert space V , either V contains a non-zero vector which is invariant under the action of G , or otherwise we have

$$\sup_{g \in K} \|\rho(g)v - v\| \geq \varepsilon \|v\|$$

for any $v \neq 0$ in V . (One says that if ρ has no invariant vector, then it also does not have *almost invariant* vectors).

If we restrict ourselves to a discrete finitely generated group G , which is the case in the applications in this book, then one shows that G has Property (T) if and only if the following holds: given an arbitrary finite generating set S , there exists $\varepsilon > 0$, depending only on S , such that for any continuous unitary representation ρ of G on a Hilbert space V , either V contains a non-zero vector which is invariant under the action of G , or otherwise we have

$$\max_{s \in S} \|\rho(s)v - v\| \geq \varepsilon \|v\|$$

for all $v \in V$. (This seems slightly stronger than the above, but see [58, Proposition 1.15].) The pair (S, ε) is called a (T) -constant for G .

In some cases, including again in this book, it is not really necessary to consider all unitary representations of G . Certain groups, the most prominent example being $SL(2, \mathbf{Z})$, fail to have Property (T), yet satisfy an analogue property for certain important subsets of representations. Lubotzky introduced a weakening of Property (T), called *Property* (τ), to deal with such situations, and we consider this briefly in Section D.3.

D.2 Properties and examples

The following are basic properties of groups with Property (T) and provide some intuition on its nature.

- An abelian group G has Property (T) if and only if G is compact (see, e.g., [58, I.2, I.5]). For instance, in \mathbf{Z} , the irreducible unitary representations are of dimension 1, and are parametrized by the unit circle in \mathbf{C} through the map which sends $t = e(\theta) \in \mathbf{C}$ to $n \mapsto t^n = e(n\theta)$; it is then intuitively clear, and easily checked, that letting $t_k \rightarrow 1$ with $t_k \neq 1$, a sequence of one-dimensional representations is obtained which has ‘almost’ invariant vectors with higher and higher precision without having invariant vectors. Taking the direct sums of those representations gives a counterexample to Property (T). This simple example illustrates another definition of Property (T) (see [58, 1.13]): there is a natural topology¹ (due to Fell) on the set \hat{G} of irreducible unitary representations of a locally compact group G , up to isomorphism, and Property (T) is equivalent with the fact that the trivial representation 1 is *isolated* in \hat{G} for this topology.

¹ Where, intuitively, representations are close if some matrix coefficient functions are close in the uniform topology on a compact subset of G .

- If G has Property (T), then so does any quotient G/H of G modulo a closed normal subgroup (indeed, representations of G/H are a subset of those of G).
- Combining the above, if G is a discrete group with Property (T), its abelianization $G/[G, G]$ is a finite group, being both discrete and compact.
- If G is a discrete group and G has Property (T), then it is finitely generated. This fact was one of the motivating consequences of Property (T) for Kazhdan. (This gives easy examples of non-abelian groups which do not have Property (T), for instance $SL(n, \mathbf{Q})$ for $n \geq 2$ with the discrete topology.)
- If G is a locally compact group having Property (T) and H is a discrete subgroup such that $H \backslash G$ carries a G -invariant probability (or finite) measure μ , i.e., such that

$$\int_{H \backslash G} f(xg) d\mu(x) = \int_{H \backslash G} f(x) d\mu(x)$$

for any integrable function $f : H \backslash G \rightarrow \mathbf{C}$ and $g \in G$, then H also has Property (T); see [58, Corollary 3.5, Corollary 4.19]. Moreover, the converse is true: if H has Property (T), then so does G . Such subgroups H are called *lattices* in G , generalizing the standard case of lattices in \mathbf{R}^n (however note that they are not always compact, in contrast with $\mathbf{R}^n/\mathbf{Z}^n$). In particular, if $H \subset G$ has finite index, then G has Property (T) if and only if H does – take the counting measure on the finite quotient.

- The groups $SL(2, \mathbf{R})$ and $SL(2, \mathbf{Z})$ do not have Property (T): indeed, it is well known that $SL(2, \mathbf{Z})$ has a finite index subgroup which is a free group, for instance the principal congruence subgroup $\Gamma(2) = \text{Ker}(SL(2, \mathbf{Z}) \rightarrow SL(2, \mathbf{Z}/2\mathbf{Z}))$, and non-trivial free groups – having infinite abelianization – do not have (T), so that the last item would bring a contradiction if $SL(2, \mathbf{Z})$ had Property (T); similarly, $SL(2, \mathbf{R}) \backslash SL(2, \mathbf{Z})$ carries a well known invariant finite measure, namely $(2\pi)^{-1} y^{-2} dx dy d\theta$ in terms of the diffeomorphism

$$\left\{ \begin{array}{l} \mathbf{R} \times]0, +\infty[\times \mathbf{S}^1 \rightarrow SL(2, \mathbf{R}) \\ (x, y, e(\theta)) \mapsto \begin{pmatrix} 1 & x \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \sqrt{y} & 0 \\ 0 & 1/\sqrt{y} \end{pmatrix} \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \end{array} \right.$$

so that the total measure is $\pi/3$ (but the quotient is not compact), and we can again apply the previous item to see that $SL(2, \mathbf{R})$ does not have Property (T).

- Note that this argument is fairly simple (including the details omitted here), compared with what would be necessary for a ‘direct’ proof by classifying the unitary irreducible representations of $SL(2, \mathbf{R})$ and checking

in the Fell topology that the trivial representation is not isolated. However, since this classification (due to Bergmann) is important in itself and provides a good comparison point, we recall that (continuous) irreducible unitary representations of $SL(2, \mathbf{R})$ can be parametrized by the union of the following four subsets of \mathbf{C} , together (in some cases) with a sign ± 1 :

- (1) The point $s = 0$, corresponding to the trivial representation (there is no sign);
- (2) The open line segment $]0, 1/2[\subset \mathbf{R}$, corresponding to the so-called ‘complementary series’ (there is no sign);
- (3) The half-vertical line $\operatorname{Re}(s) = 1/2$ with $\operatorname{Im}(s) > 0$, corresponding to the ‘principal series’, with a sign ± 1 , except for $s = 1/2$ where there is only one sign;
- (4) The points $s = k(1 - k/2)/2$ for $k \geq 2$ an integer, corresponding to the ‘discrete series’, with sign ± 1 , and the ‘limit of discrete series’ for $s = 1/4$, with no sign (note that $s = 1/4$ is a ‘double point’, arising also as a principal series).

(See, e.g., [79, Theorem 16.3] for this statement, with different normalizations.) Except for $s = 0$, all these representations are infinite-dimensional – this goes a long way towards explaining why this classification is highly non-trivial. Of course, the same is not true for $SL(2, \mathbf{Z})$, which has plenty of finite-dimensional irreducible representations factoring through a finite quotient such as $SL(2, \mathbf{Z}/n\mathbf{Z})$.

In concrete terms, the parameter $s \in \mathbf{C}$ has the following interpretation: the eigenvalue of the so-called Casimir operator (properly normalized) acting on the representation is equal to $s(1 - s)$. Except for discrete series, note that it is a non-negative real number.

If we grant the fact (also by no means obvious) that the Fell topology on the set of irreducible representations is the same as that induced by seeing the set of parameters s as a subset of \mathbf{C} , then we see that the reason that $SL(2, \mathbf{R})$ does not have Property (T) is that the point $s = 0$ (representing the trivial representation) can be approached by means of the complementary series of representations.

These facts have the following interpretation in the theory of automorphic forms: to any Maass (non-holomorphic) cusp form $f \in L^2(\Gamma \backslash \mathbf{H})$ on the Poincaré upper half-plane \mathbf{H} with respect to a discrete subgroup Γ of $SL(2, \mathbf{R})$ with finite covolume, there is associated (in a natural way) a representation of $SL(2, \mathbf{R})$ with parameter s such that $\lambda = s(1 - s)$, where λ is the Laplace eigenvalue of f , i.e.,

$$-y^2 \left(\frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2} \right) = \lambda f.$$

The eigenvalue is then $\geq 1/4$ if and only if the representation belongs to the principal series, and $0 < \lambda < 1/4$ (i.e., this is an exceptional eigenvalue, in the usual terminology) if and only if the representation belongs to the complementary series. The well-known conjecture of Selberg on the first eigenvalue of the Laplace operator for *congruence subgroups* is then equivalent with the conjecture that all cusp forms for such a group are always associated with a principal series representation. The theorem of Selberg according to which $\lambda \geq 3/16$ for congruence subgroups means that, although complementary series can conceivably occur, they can not be ‘too close’ to the trivial representation – hence the latter is isolated *among automorphic representations associated to congruence subgroups*, and this is precisely the crucial result required to prove that $SL(2, \mathbf{Z})$ has Property (τ) with respect to the family of such subgroups. (See the discussion in Section 7.4 for references.)

- Now for examples of (non-compact) groups with Property (T). First of all, $SL(n, \mathbf{R})$, for $n \geq 3$, has Property (T) (see, e.g., [58, Theorem 2.4]), which is due to Kazhdan. Hence any discrete subgroup G in $SL(n, \mathbf{R})$ such that $H \backslash SL(n, \mathbf{R})$ carries a finite volume invariant measure (i.e., any lattice in $SL(n, \mathbf{R})$) also has Property (T). It is well known that $SL(3, \mathbf{Z})$ (and its finite index subgroups) have this property. Alternatively, as already mentioned, Shalom proved directly that $SL(n, \mathbf{Z})$ ($n \geq 3$) has Property (T), and we will sketch his proof below.
- In addition, the groups $Sp(2g, \mathbf{R})$ for $g \geq 2$ have Property (T), hence so do the lattices in G , among them $Sp(2g, \mathbf{Z})$ and its finite index subgroups. Shalom’s method has been extended to this case by Neuhauser [101].
- Very recently, Shalom has also proved that $SL(n, \mathbf{Z}[x_1, \dots, x_m])$ has Property (T) for all $n \geq m + 3$ (the weaker Property (τ) had been proved earlier by Kassabov and Nikolov), see [119] for a sketch of the proof.

D.3 Property (τ)

As explained in the previous section, neither $SL(2, \mathbf{R})$ nor $SL(2, \mathbf{Z})$ have Property (T). However, motivated by the fact that it was known that certain important *subsets* of unitary representations still satisfied the defining separation condition, Lubotzky introduced Property (τ) as a weakening of Property (T). See [91] and [93] for more detailed discussions.

Let G be a finitely generated group, and let $(N_i)_{i \in I}$ be an arbitrary family of finite index subgroups of G . The group G is said to have Property (τ) with

respect to (N_i) if there exist a finite set S and $\varepsilon > 0$ such that for any unitary representation ρ of G on a Hilbert space V which satisfies $\text{Ker } \rho \supset N_i$ for some i and does not contain a non-zero invariant vector, we have

$$\max_{s \in S} \|\rho(s)v - v\| \geq \varepsilon \|v\|$$

for all non-zero $v \in V$. In this situation, the pair (S, ε) is called a (τ) -constant for $(G, (N_i))$. If the family (N_i) is not mentioned, it is implicitly taken to be the family of all finite-index subgroups.

The basic example is that $SL(2, \mathbf{Z})$ has Property (τ) with respect to the set of congruence subgroups $\Gamma(d) = \text{Ker}(SL(2, \mathbf{Z}) \rightarrow SL(2, \mathbf{Z}/d\mathbf{Z}))$; see Section 7.4 for references concerning this fact, which is essentially a version of Selberg's theorem about the first eigenvalue of the hyperbolic Laplacian acting on $L^2(\Gamma(d) \backslash \mathbf{H})$. Generalizing this, the group $SL(2, \mathbf{Z}_K)$, where \mathbf{Z}_K is the ring of integers in a number field K , has Property (τ) with respect to congruence subgroups $\text{Ker}(SL(2, \mathbf{Z}_K) \rightarrow SL(2, \mathbf{Z}_K/I))$ for all ideals $I \subset \mathbf{Z}_K$.

Property (τ) does not have quite the same stability properties as Property (T) does: for instance, $\Gamma(2)$ has Property (τ) with respect to the congruence subgroups $\Gamma(2d)$, $d \geq 1$, yet its abelianization is infinite. However, Property (τ) with respect to all finite-index subgroups (which is the meaning of Property (τ) when no family of subgroups is indicated explicitly) does imply that the abelianization is finite. In another direction, there exist groups G containing lattices G_1 and G_2 , yet G_1 has Property (τ) whereas G_2 does not.

Generalizing Selberg's theorem, a result of Clozel [23] shows that for any simply-connected semisimple algebraic group G over a number field k with ring of integers \mathbf{Z}_k , the group $G(\mathbf{Z}_k)$ has Property (τ) with respect to the family of congruence subgroups

$$\text{Ker}(G(\mathbf{Z}_k) \rightarrow G(\mathbf{Z}_k/I))$$

where $I \subset \mathbf{Z}_k$ ranges over non-zero ideals.

There is a far-reaching conjecture that states (over \mathbf{Q}) that for such a group G , for any Zariski-dense discrete subgroup $\Gamma \subset G(\mathbf{Z})$, Γ should have Property (τ) with respect to the family of 'congruence' subgroups

$$\text{Ker}(\Gamma \rightarrow G(\mathbf{Z}/q\mathbf{Z}))$$

for q squarefree (see [14, Conjecture 1.4]). Due to the recent results of Helfgott [59], Bourgain and Gamburd [13] and Bourgain, Gamburd and Sarnak [15], this conjecture is now known for the case of $SL(2)$.

D.4 Shalom's theorem

Shalom [118] gave the first proof of Property (T) for $SL(n, \mathbf{Z})$ that did not involve seeing it as a subgroup of $SL(n, \mathbf{R})$. His proof is quite short and its only ‘technical’ tool (the harmonic analysis of infinite dimensional representations of abelian groups) can be further eliminated if only finite-dimensional representations are considered. In fact, in this case the result was proved earlier (for $SL(3, \mathbf{Z})$ at least) by M. Burger [18], whose ideas form an important part of Shalom's argument. In particular, Property (τ) with respect to all finite-index subgroups follows in this manner. Since this is exactly what is needed in Chapter 7, we include a fairly detailed sketch. As in [118], the only ingredient we quote from the literature is the ‘bounded elementary generation property’ of $SL(n, \mathbf{Z})$; for this, and for another complete exposition, see [9, Chapter 4]. In fact, as pointed out by M. Burger, his own argument leads to the same result while avoiding the bounded generation property (and, although written for $SL(3, \mathbf{Z})$, it can be extended to $SL(n, \mathbf{Z})$), with a better constant; however it is rather lengthier ([18, Section 3]).

Theorem D.1 (Shalom) *Let $n \geq 3$ be an integer and $S = S^{-1}$ the symmetric generating set of $SL(n, \mathbf{Z})$ consisting of elementary matrices with ± 1 off the diagonal. Then for any continuous finite-dimensional unitary representation*

$$\rho : SL(n, \mathbf{Z}) \rightarrow U(V)$$

where V is a finite-dimensional Hilbert space, one of the following holds:

- *there is a non-zero vector $v \in V$ invariant under $SL(n, \mathbf{Z})$;*
- *for any $v \neq 0$ in V , there exists $s \in S$ such that $\|\rho(s)v - v\| \geq \varepsilon_n v$, where*

$$\varepsilon_n = \frac{1}{(4 + \sqrt{21})(3n^2 - n + 102)}.$$

The first step in the proof (which is related to the earliest proofs of Property (T), see, e.g., [58, Chapter 2], and [18, Section 1, Section 5]) is to look at the restriction of the representations to subgroups of a special type.

Lemma D.2 (Shalom) *Let V be a finite-dimensional Hilbert space and let π be a continuous unitary representation of the semi-direct product*

$$G = \mathbf{Z}^2 \rtimes SL(2, \mathbf{Z})$$

on V . Let $F \subset G$ be the set

$$F = \left\{ ((\pm 1, 0), \text{Id}), ((0, \pm 1), \text{Id}), (0, \begin{pmatrix} 1 & \pm 1 \\ 0 & 1 \end{pmatrix}), (0, \begin{pmatrix} 1 & 0 \\ \pm 1 & 1 \end{pmatrix}) \right\}.$$

- (1) If no non-zero vector in V is invariant under the action of the subgroup $\mathbf{Z}^2 \subset G$, then for any non-zero unit vector $v \in V$, there exists an element $s \in F$ such that

$$\|\pi(s)v - v\| \geq \frac{-4 + \sqrt{21}}{5} \simeq 0.1165151 \dots \quad (\text{D.1})$$

- (2) With no assumptions, for any $\varepsilon > 0$, if $v \in V$ is a unit vector such that

$$\max_{s \in F} \|\pi(s)v - v\| \leq \varepsilon,$$

then we have

$$\max_{m \in \mathbf{Z}^2} \|\pi(m)v - v\| \leq (8 + 2\sqrt{21})\varepsilon.$$

Property (1) is called the *relative Property (T)* of the pair (G, \mathbf{Z}^2) (or rather relative Property (τ) here, since we are not dealing with all representations of G).

Proof Recall that G is the group with underlying set $\mathbf{Z}^2 \times SL(2, \mathbf{Z})$ and group law given by

$$(m, g) \cdot (n, h) = (m + g \cdot n, gh), \quad (m, g)^{-1} = (-g^{-1} \cdot m, g^{-1})$$

where $SL(2, \mathbf{Z})$ acts on \mathbf{Z}^2 , seen as column vectors, in the standard way. The subgroup \mathbf{Z}^2 is normal in G , and the quotient G/\mathbf{Z}^2 is isomorphic to $SL(2, \mathbf{Z})$.

(1) This is the most crucial part of the proof. We reproduce Shalom's argument, except for 'running it backwards' (which we do simply in order to avoid a simple copy of his own very clear writing). In doing so, we obtain a slightly better relative Kazhdan constant,² but this improvement is of course immaterial.

Restricting the representation π to the subgroup $\mathbf{Z}^2 \subset G$, we can use the representation theory of abelian groups to obtain an orthogonal decomposition³

$$V = \bigoplus_{\xi \in T} V_{\xi} \quad (\text{D.2})$$

where ξ runs over a finite subset $T \subset (\mathbf{R}/\mathbf{Z})^2$ and V_{ξ} is the space of vectors $v \in V$ such that

$$\pi(m)v = e(\langle m, \xi \rangle)v, \quad \langle m, \xi \rangle = m_1\xi_1 + m_2\xi_2,$$

for all $m = (m_1, m_2) \in \mathbf{Z}^2$ and $\xi = (\xi_1, \xi_2) \in T$. There exists $v \neq 0$ in V which is \mathbf{Z}^2 -invariant if and only if the trivial character occurs in this decomposition, namely if $0 \in T$. Moreover, because the representation π is a representation

² Shalom has $1/10$ instead of our $0.11 \dots$

³ This is where it simplifies matters to have a finite-dimensional representation V .

of G , and \mathbf{Z}^2 is normal in G , a simple computation shows that for any $g \in SL(2, \mathbf{Z})$, we have an isometry

$$\begin{cases} V_\xi \rightarrow V_{g \cdot \xi} \\ v \mapsto \pi(g)v \end{cases}$$

where $g \cdot \xi$ denotes the left action of $SL(2, \mathbf{Z})$ on $(\mathbf{R}/\mathbf{Z})^2$ (seen as column vectors) by the ‘standard’ (matrix product) action of the inverse-transpose of g .

Now assume there is no non-zero vector invariant under \mathbf{Z}^2 , and fix a unit vector $v \neq 0$. We define

$$\varepsilon = \max_{s \in F \cap \mathbf{Z}^2} \|\pi(s)v - v\| > 0$$

since $F \cap \mathbf{Z}^2$ generates \mathbf{Z}^2 . Let $\delta, 0 < \delta < 1$, be given, and put

$$T_\delta = \{\xi \in T \mid \|\xi\| < \delta\} \subset T$$

where $\|\xi\| = \max(\|\xi_1\|, \|\xi_2\|)$ with $\|\cdot\|$ denoting the distance to 0 in \mathbf{R}/\mathbf{Z} . Correspondingly, with $v_\xi \in V_\xi$ the ξ -component of the vector v , let

$$w = \sum_{\xi \in T_\delta} v_\xi.$$

For $m = (\pm 1, 0)$ or $(0, \pm 1)$ (in $F \cap \mathbf{Z}^2$), we have (with obvious notation) the relation

$$\|\pi(m)v - v\|^2 = \sum_{\xi \in T} |e(\langle m, \xi \rangle) - 1|^2 \|v_\xi\|^2,$$

and it follows by combining these four equalities with the definition of T_δ that we have

$$\|w\|^2 \geq 1 - \frac{1}{2} \left(\frac{\varepsilon}{\sin \pi \delta} \right)^2$$

(thus, if v is almost invariant under $F \cap \mathbf{Z}^2$, so that ε is small, ‘most’ of the vector is supported on V_ξ with ξ close to the trivial character, which is rather intuitive).

Next, following an idea going back to Burger, we partition $(\mathbf{R}/\mathbf{Z})^2 - \{0\}$, identified with $] -1/2, 1/2]^2 - \{0\}$, into four subsets A, B, C and D , as described in Figure D.1 (where the inner square represents the boundary of T_δ).⁴

Since $0 \notin T$, one of them, say $Y \in \{A, B, C, D\}$, is such that

$$\sum_{\xi \in Y \cap T(\delta)} \|v_\xi\|^2 \geq \frac{1}{4} \|w\|^2. \quad (\text{D.3})$$

⁴ The boundaries being half-open, half-closed in a clockwise direction, say.




Figure D.1 The four regions of the Burger lemma

Assume that $Y = B$ for instance. Then, letting $X = A \cup B$ and

$$s = \begin{pmatrix} 1 & 0 \\ -1 & 1 \end{pmatrix},$$

notice that (with the inverse-transpose action on $(\mathbf{R}/\mathbf{Z})^2$) we have $s \cdot X = A$ if $\delta \leq 1/4$, this condition ensuring that the action of s restricted to T_δ is injective. Hence we have a disjoint union

$$X \cap T_\delta = (s \cdot X \cap T_\delta) \cup (B \cap T_\delta),$$

and it follows from the inequality (D.3) that

$$\mu_\delta(X) - \mu_\delta(s \cdot X) \geq \frac{\|w\|^2}{4} \geq \frac{1}{4} \left(1 - \frac{1}{2} \left(\frac{\varepsilon}{\sin \pi \delta} \right)^2 \right),$$

where $\mu_\delta(Y)$, for any subset Y , is defined as the sum of the $\|v_\xi\|^2$ for those $\xi \in Y \cap T_\delta$ (recall that the vector v is fixed).

Now write $\mu(Y)$ for the sum of the squared norms of v_ξ for all $\xi \in Y$. We have

$$\begin{aligned} \mu(X) - \mu(s \cdot X) &= \mu_\delta(X) - \mu_\delta(s \cdot X) - (\mu_\delta(X) - \mu(X)) \\ &\quad - (\mu(s \cdot X) - \mu_\delta(s \cdot X)) \\ &\geq \frac{1}{4} \left(1 - \frac{1}{2} \left(\frac{\varepsilon}{\sin \pi \delta} \right)^2 \right) - \frac{1}{2} \left(\frac{\varepsilon}{\sin \pi \delta} \right)^2 \\ &= \frac{1}{4} \left(1 - \frac{5}{2} \left(\frac{\varepsilon}{\sin \pi \delta} \right)^2 \right) \end{aligned}$$

since $0 \leq \mu(X) - \mu_\delta(X) \leq \|v - w\|^2 = 1 - \|w\|^2$. Finally, observe that, on the other hand, we have

$$\begin{aligned} \mu(X) - \mu(s \cdot X) &= \langle Pv, v \rangle - \langle \pi(s^{-1})P\pi(s)v, v \rangle \\ &= \langle Pv, v - \pi(s)v \rangle - \langle \pi(s^{-1})P(\pi(s)v - v), v \rangle \\ &\leq 2\|\pi(s)v - v\|, \end{aligned}$$

where P is the orthogonal projection from V to the sum of the V_ξ with $\xi \in X$. From this it follows that

$$\|\pi(s)v - v\| \geq \frac{1}{2}(\mu(X) - \mu(s \cdot X)) \geq \frac{1}{8} \left(1 - \frac{5}{2} \left(\frac{\varepsilon}{\sin \pi \delta} \right)^2 \right),$$

and altogether we derive

$$\max_{s \in F} \|\pi(s)v - v\| \geq \max \left(\varepsilon, \frac{1}{8} \left(1 - \frac{5}{2} \left(\frac{\varepsilon}{\sin \pi \delta} \right)^2 \right) \right).$$

This inequality is obviously best for δ as large as possible, and taking $\delta = 1/4$, it becomes

$$\max_{s \in F} \|\pi(s)v - v\| \geq \max \left(\varepsilon, \frac{1 - 5\varepsilon^2}{8} \right),$$

and finally this function of $\varepsilon > 0$ achieves its lower bound for $\varepsilon = (-4 + \sqrt{21})/5$, so that we obtain (D.1), finishing the proof of (1).

(2) We can decompose V as an orthogonal direct sum $V = V_0 \oplus V_1$ where V_0 is the space of vectors invariant under the action of $\mathbf{Z}^2 \subset G$ and V_1 its orthogonal complement. Those subspaces are clearly invariant under \mathbf{Z}^2 , and because \mathbf{Z}^2 is a normal subgroup, they are in fact G -invariant. Writing $v = v_0 + v_1$ for a unit vector v with $v_i \in V_i$, we have for $s \in F$ the bound

$$\|\pi(s)v - v\|^2 = \|\pi(s)v_0 - v_0\|^2 + \|\pi(s)v_1 - v_1\|^2 \leq \varepsilon^2.$$

Let $\varepsilon_0 = (-4 + \sqrt{21})/5$. By (1), applied to the representation of G on V_1 , we can select one $s \in F$ such that $\|\pi(s)v_1 - v_1\|^2 \geq \varepsilon_0^2 \|v_1\|^2$, and this leads to $\|v_1\| \leq \varepsilon_0^{-1} \varepsilon$. Now, turning back to an arbitrary $m \in \mathbf{Z}^2$, we also have

$$\begin{aligned} \|\pi(m)v - v\|^2 &= \|\pi(m)v_0 - v_0\|^2 + \|\pi(m)v_1 - v_1\|^2 \\ &= \|\pi(m)v_1 - v_1\|^2 \leq 4\|v_1\|^2 \leq 4\varepsilon_0^{-2} \varepsilon^2, \end{aligned}$$

hence the result. \square

With this lemma in hand, the proof of Theorem D.1 proceeds as follows: assume that there is no non-zero invariant vector under $SL(n, \mathbf{Z})$ in V , and then let $v \in V$ be a unit vector and define

$$\varepsilon = \varepsilon(v) = \max_{s \in S} \|\rho(s)v - v\| > 0.$$

For any elementary matrix g (i.e., of the form s^m for some $s \in S$ and $m \geq 1$), it is possible to find⁵ a subgroup G of $SL(n, \mathbf{Z})$ containing g such that $G \simeq \mathbf{Z}^2 \rtimes SL(2, \mathbf{Z})$ and g is in the subgroup \mathbf{Z}^2 with this identification (see Lemma 2.4 in [118]; e.g., if $n = 3$ and

$$g = \begin{pmatrix} 1 & 0 & 12 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix},$$

one can take for G the set of matrices of the type

$$\begin{pmatrix} a & b & x \\ c & d & y \\ 0 & 0 & 1 \end{pmatrix},$$

with $ad - bc = 1$, giving the $SL(2, \mathbf{Z})$ part, and $(x, y) \in \mathbf{Z}^2$; the reader is encouraged to check that the group law corresponds to the one of the semi-direct product). Moreover, this group G is such that $S \cap G$ corresponds to the set F of Lemma D.2, and hence according to part (2) of this lemma (applied with π given by the restriction of ρ to G), it follows from the definition of ε that

$$\|\rho(g)v - v\| \leq (8 + 2\sqrt{21})\varepsilon. \quad (\text{D.4})$$

Now let $h \in SL(n, \mathbf{Z})$ be arbitrary. The next crucial property is the following: we can write

$$h = g_1 \cdots g_k$$

with g_i elementary, and k bounded independently of h . This further property of $SL(n, \mathbf{Z})$ and S is called the *bounded elementary generation property*, and it is known to hold with $k = \frac{1}{2}(3n^2 - n) + 51$, as shown by Carter and Keller (see the references and discussion in [118] for a more general discussion, or [9, Section 4.1]; the proof of this result depends on Dirichlet's theorem on primes in arithmetic progressions).

By splitting

$$\begin{aligned} \|\rho(h)v - v\| &\leq \sum_{0 \leq i \leq k-1} \|\rho(g_0 \cdots g_{k-i})v - \rho(g_0 \cdots g_{k-i-1})v\| \\ &= \sum_{1 \leq j \leq k} \|\rho(g_j)v - v\|, \end{aligned}$$

using unitarity, we obtain by (D.4) the bound

$$\|\rho(h)v - v\| \leq (8 + 2\sqrt{21})\varepsilon k.$$

⁵ This uses that $n \geq 3$.

Here is now the final flourish: if

$$(8 + 2\sqrt{21})\varepsilon k \leq 1 \tag{D.5}$$

then we have found that $\|\rho(h)v - v\| \leq 1$ for all $h \in SL(n, \mathbf{Z})$. But then a well-known fact about Hilbert spaces shows that since the orbit $Q = \rho(SL(n, \mathbf{Z}))v$ of v in V is bounded, there is a unique centre of mass $w \in V$ that minimizes the distance to Q . By uniqueness, this w is invariant under $SL(n, \mathbf{Z})$, and the bound ≤ 1 ensures that $w \neq 0$ (recall v is a unit vector, so $w = 0$ would contradict the uniqueness of the centre of mass). So the inequality (D.5) is impossible by the assumption that V has no (non-zero) invariant vector, and hence

$$\varepsilon > \frac{1}{(8 + 2\sqrt{21})k},$$

which precisely gives Property (T) (or rather, (τ), since we only consider finite-dimensional representations), with the Kazhdan constant that was claimed.

Remark D.3 To prove the full Property (T) for $SL(n, \mathbf{Z})$, what is needed is the spectral decomposition theory for general (possibly infinite-dimensional) representations of abelian groups, specifically for \mathbf{Z}^2 (to generalize Lemma D.2). In this generality, the orthogonal direct sum decomposition (D.2) is replaced by a ‘direct integral’ over the whole set $(\mathbf{R}/\mathbf{Z})^2$ of characters of \mathbf{Z}^2 . What this essentially means is that, for a given unit vector v , there exists a probability measure μ_v on the Borel subsets of $(\mathbf{R}/\mathbf{Z})^2$ such that $\mu_v(X)$ is (intuitively) the squared norm of the projection of the vector v onto the space spanned by the components of the direct integral parametrized by $\xi \in X \subset (\mathbf{R}/\mathbf{Z})^2$. For a finite-dimensional representation decomposed as in (D.2), this measure is a sum of Dirac measures at $\xi \in X$ weighted by $\|v_\xi\|^2$. Shalom’s paper describes very clearly how this works.

Appendix E

Linear algebraic groups

This chapter is simply a list of basic definitions related to the theory of linear algebraic groups, included for completeness so that readers unfamiliar with this language can understand the statements and proofs in Chapters 5 and 7.

E.1 Basic terminology

We use the language of varieties, identifying an algebraic variety with the set of its points over an algebraically closed field; this is indeed sufficient for much of the theory of algebraic groups. For references, see for instance the books of Borel [12] or Springer [125]. Note that there are subtle issues of regularity and rationality involved when the base field is not perfect (see the examples in [125, 12.1.6]), and we will therefore assume that K is perfect (the cases of interest to us being K finite or of characteristic zero).

- A *linear algebraic group* defined over a perfect subfield K of an algebraically closed field \overline{K} is a subgroup $G \subset GL(n, \overline{K})$ for some $n \geq 1$ defined by a set of polynomial equations involving the coordinates of a matrix and the inverse of its determinant: there exist polynomials f_1, f_2, \dots , with coefficients in K , in $n^2 + 1$ variables, such that $g \in G$ if and only if

$$f_1(g_{i,j}, (\det(g_{i,j}))^{-1}) = \dots = f_m(g_{i,j}, (\det(g_{i,j}))^{-1}) = \dots = 0$$

for all m . Even if the set of equations is infinite, Hilbert's theorem (polynomial rings in finitely many variables over a field are noetherian) implies that one can reduce to a finite set of equations, namely generators of the ideal generated by all polynomials.

- For any field L such that $K \subset L$, the set of solutions of those equations in L is denoted $G(L)$, and it is a subgroup of $GL(n, L)$. The group G , when seen as defined over L , is denoted G/L .

- A homomorphism of linear algebraic groups G and H defined over K is a group homomorphism $G \xrightarrow{\varphi} H$ such that each coordinate of $\varphi(g)$ is given by polynomials in the same $n^2 + 1$ variables as above. The kernel of such a homomorphism is then clearly a linear algebraic group, and so is the image $\varphi(G)$, less obviously so. If G is defined over K , $\text{Ker } \varphi$ and $\text{Im } \varphi$ are also defined over K .
- Products of linear algebraic groups are defined in the obvious manner. The centre $Z(G)$ of a linear algebraic group G is one itself (using first infinitely many equations to describe the intersection of the centralizers of all elements of G), and is defined over K if G is.
- The basic examples are $GL(n, \overline{K})$, with $GL(1, \overline{K})$ called the *multiplicative group*, often denoted simply by $GL(n)/K$ or \mathbf{G}_m/K . The determinant map $\det : GL(n) \rightarrow \mathbf{G}_m$ is a homomorphism of linear algebraic groups, and its kernel $SL(n)/K$ is therefore a linear algebraic group. Also we have the *additive group*, isomorphic to \overline{K} with the addition law, denoted by \mathbf{G}_a/K ; it may be identified with the subgroup of $GL(2, \overline{K})$ given by

$$\left\{ \begin{pmatrix} 1 & x \\ 0 & 1 \end{pmatrix} \mid x \in \overline{K} \right\},$$

for instance.

- If K^{2g} , for some $g \geq 1$, is equipped with a non-degenerate alternating bilinear form $\langle \cdot, \cdot \rangle$, the *symplectic group* $Sp(2g)$, defined as the subgroup of $GL(2g)$ of matrices leaving the bilinear form invariant, i.e.,

$$Sp(2g) = Sp(\langle \cdot, \cdot \rangle) = \{g \in GL(2g, \overline{K}) \mid \langle gv, gw \rangle = \langle v, w \rangle \\ \text{for all } v, w \in \overline{K}^{2g}\}$$

is clearly a linear algebraic group over K . Similarly, the *group of symplectic similitudes* $CSp(2g)$ defined by

$$CSp(2g) = CSp(\langle \cdot, \cdot \rangle) = \{g \in GL(2g, \overline{K}) \mid \langle gv, gw \rangle = \lambda \langle v, w \rangle \\ \text{for all } v, w \in \overline{K}^{2g} \text{ and some } \lambda \in \overline{K}\}$$

is a linear algebraic group over K . The multiplier map $m : CSp(2g) \rightarrow \mathbf{G}_m$ mapping g to λ is a homomorphism of linear algebraic groups.

- A linear algebraic group G over K is *connected* if there is no proper finite index subgroup H which is itself a linear algebraic group over K . It is *geometrically connected*, if there is no proper finite index subgroup which is a linear algebraic group over \overline{K} . There exists a maximal connected subgroup of

G , which is normal and of finite index in G , and is called the *connected component of the identity* of G . The groups $GL(n)$, $SL(n)$ and \mathbf{G}_a are connected, and so are $Sp(2g)$ and $CSp(2g)$.

- A *torus* defined over K is a linear algebraic group which is isomorphic over \overline{K} to a product of copies of the multiplicative group \mathbf{G}_m ; if the isomorphism is defined over K itself, the torus is said to be *split*. For example, if the characteristic of K is odd, defining

$$G = \{(x, y) \in \overline{K} \mid x^2 + y^2 = 1\}$$

with group law

$$(x, y) \cdot (x', y') = (xx' - yy', xy' + x'y),$$

it is easy to check that G is a linear algebraic group, seen as a subgroup of $GL(2)$ by means of the map

$$(x, y) \mapsto \begin{pmatrix} x & -y \\ y & x \end{pmatrix}.$$

Over \overline{K} (indeed, over any field containing a square root i of -1), this group is isomorphic to \mathbf{G}_m^2 by the map

$$(x, y) \mapsto (x + iy, x - iy), \text{ with inverse } (z, w) \mapsto (\tfrac{1}{2}(z + w), \tfrac{1}{2i}(z - w)),$$

but G is not isomorphic to \mathbf{G}_m^2 over K itself if $i \notin K$, so in that case it is a non-split torus. As another example, the centre of $GL(n)$ is a torus isomorphic to \mathbf{G}_m , and the group of diagonal matrices is a split torus in $GL(n)$ isomorphic to \mathbf{G}_m^n .

- For any field L containing K , the L -rank of G is defined to be the greatest integer $r \geq 0$ such that L contains a torus which is split over L and isomorphic (over L) to \mathbf{G}_m^r . If L is algebraically closed, the L -rank is called the rank of G , and any torus in G which is of dimension equal to the rank is called a *maximal torus*. All maximal tori are conjugate in G . For instance, $GL(n)$ is of rank n and $SL(n)$ is of rank $n - 1$ (over any L). The non-split torus above is of K -rank 1 over a field not containing i , and of rank 2 over $K(\sqrt{-1})$. If the K -rank of G is equal to the rank of G , i.e., if there exists a maximal torus defined over K , then G is said to be *split*. In this case all K -rational maximal tori are K -conjugate. This is the case of $GL(n)$, $SL(n)$, $Sp(2g)$, $CSp(2g)$.
- A unipotent subgroup of G is any subgroup containing only unipotent elements, i.e., matrices g such that $(g - \text{Id})^m = 0$ for some $m \geq 1$. Although this seems to depend on the choice of an embedding in a matrix group, this condition is independent of such a choice. In particular, \mathbf{G}_a is unipotent. In $GL(n)$, the subgroup U of upper-triangular matrices with 1 on the diagonal

is a unipotent subgroup, and in fact it is a maximal unipotent subgroup, and any unipotent subgroup is conjugate to a subgroup of U .

- An element $g \in G$ is *semisimple* if it is contained in a torus in G , or equivalently if $G \subset GL(n)$, if $g \in GL(n)$ is diagonalizable. Moreover, g is regular if it is contained in a unique maximal torus. In $GL(n)$, this means g has distinct eigenvalues. Note that those definitions are not the most standard ones. Any element $g \in G$ can be written uniquely $g = g_s g_u$ where $g_s \in G$ is semisimple, $g_u \in G$ is unipotent, and g_s, g_u commute. If G is defined over K , $g \in G(K)$, then g_s and g_u are also in $G(K)$ (this is one place where having K perfect matters).
- A *Borel subgroup* in a connected linear algebraic group G is a maximal connected solvable subgroup. All such subgroups are conjugate. In the case of $GL(n)$, the standard example is the subgroup B of upper-triangular matrices. In particular, the following general facts are clear in that case: a Borel subgroup B contains a maximal torus T (the diagonal matrices, in the example) which is the centre of B and a maximal unipotent subgroup U of G ; moreover, B is the semi-direct product of T and U .
- A linear algebraic group G is *reductive* if and only if it is connected¹ and contains no non-trivial connected normal unipotent subgroup. It is *semisimple* if it is connected and contains no non-trivial connected normal abelian subgroup.² A semisimple group is also reductive. Another characterization is the following: G is reductive if and only if there exists a representation $G \rightarrow GL(n)$ with finite kernel which is completely reducible, i.e., a direct sum of irreducible representations.
- An important property of reductive groups is that each maximal torus is its own centralizer. (In general, the centralizer of a maximal torus is called a *Cartan subgroup*, so for a reductive group, Cartan subgroups and maximal tori coincide.) A reductive group, or more generally a connected linear algebraic group, G , is *simply-connected* if any surjective morphism $H \rightarrow G$ with finite kernel where H is a connected linear algebraic group is an isomorphism. (Be careful with the relation with the usual topological definition recalled in Appendix H: for instance $SL(2, \mathbf{R})$ has fundamental group isomorphic to \mathbf{Z} , whereas $SL(2)$ is simply-connected as an algebraic group; for the group of complex points, e.g., $SL(2, \mathbf{C})$, there is no problem.) For instance, $PSL(n)$, the quotient of $SL(n)$ modulo its (finite) centre, which is the group of scalar matrices with n -th roots of unity, is not simply-connected, as shown by the surjective map $SL(n) \rightarrow PSL(n)$.

¹ This condition is sometimes omitted.

² Note that a semisimple group is *not* a group where all elements are semisimple

- The group $GL(n)$ itself is reductive (this is obvious from the ‘completely reducible representation’ point of view); its centre is isomorphic to \mathbf{G}_m , showing that $GL(n)$ is not semisimple. However, $SL(n)$ is semisimple. Similarly, the group $CSp(2g)$ is reductive (again, the representation $CSp(2g) \rightarrow GL(2g)$ is faithful and irreducible) and not semisimple, while $Sp(2g)$ is semisimple. These four groups are also simply-connected. On the other hand, a non-trivial unipotent group is not reductive (by the first definition).
- More generally, if G is reductive, then $[G, G]$ is a linear algebraic group which is semisimple, and moreover $G = [G, G]Z(G)$, with the intersection $Z(G) \cap [G, G]$ being finite. The rank of $[G, G]$ is called the *semisimple rank* of G .
- If G is a reductive linear algebraic group, and T is a maximal torus, the *Weyl group* of G is defined as the quotient $N(T)/T$ where T is the normalizer of T in G . This turns out to be a finite group, and (up to isomorphism) it does not depend on the choice of T . If $G = GL(n)$ or $SL(n)$, T may be chosen to be the group of diagonal matrices, and then $N(T)$ is the semi-direct product of T and the finite group of permutation matrices, so that W is isomorphic to the symmetric group \mathfrak{S}_n . If $G = Sp(2g)$ or $CSp(2g)$, one finds that W is isomorphic to the group W_{2g} of permutations σ of $\{1, \dots, 2g\}$ which act on pairs $\{2i-1, 2i\}$, $1 \leq i \leq g$; in other words this is the same group as the one which occurs prominently in Chavdarov’s problem in Chapter 8.

E.2 Galois groups of characteristic polynomials

As an example of use of the previous notions, we sketch a proof of the fact that the Galois group of the splitting field of the characteristic polynomial of a matrix $g \in SL(n, \mathbf{Q})$, $GL(n, \mathbf{Q})$, $Sp(2g, \mathbf{Q})$ or $CSp(2g, \mathbf{Q})$ is isomorphic to a subgroup of the Weyl group of the corresponding algebraic group $SL(n)$, $GL(n)$, $Sp(2g)$, $CSp(2g)$, which we checked ‘by hand’ at the beginning of the proof of Theorem 7.12 and in Section E.1.

Proposition E.1 *Let K be a field, and let G be one of $GL(n)$, $SL(n)$, $Sp(2g)$, $CSp(2g)$ for some $n \geq 1$ or $g \geq 1$, or a product of such groups.*

Let $g \in G(K)$ be a regular semisimple element and let L/K be the splitting field of $\det(X - g) \in K[X]$.³ Then there is an injective homomorphism $\text{Gal}(L/K) \rightarrow W$.

³ If G is a product of groups G_1, \dots, G_k , of the type described, $G_i \subset GL(d_i)$, the characteristic polynomial is computed in $GL(d_1 + \dots + d_k)$.

Proof The assumptions on G which will really be used are that G is (or is K -isomorphic to) a split, simply-connected, connected, reductive algebraic group over K , which is a subgroup of $GL(d)$ for some $d \geq 1$, in such a way that G has a K -maximal torus T which is a subgroup of the group of diagonal matrices in $GL(d)$. The characteristic polynomial is then computed in $GL(d)$.

Now, consider the set

$$X_g = \{t \in T \mid t \text{ and } g \text{ are conjugate}\}$$

(where conjugation is in G , i.e., over an algebraically closed field). We claim that the following properties hold:

- (i) The set X_g is non-empty; since g is semisimple, with our definition this follows from the fact that any maximal torus containing g is conjugate to T (see, e.g., [12, II, Theorem 11.10]).
- (ii) The Weyl group W acts naturally on X_g by conjugation; indeed, it is clear that the normalizer $N(T)$ acts by conjugation on X_g , and that the centralizer $C(T)$ acts trivially. Since G is reductive, we have $C(T) = T$ (see, e.g., [12, II.13.17, Corollary 2]), hence the required action of $W = N(T)/T$ on X_g . Note that representatives of W in $N(T)$ can be chosen in $N(T)(K)$ because G is split over K (see, e.g., [125, Paragraph before 16.1.3]).
- (iii) The action of W on X_g is transitive and free, i.e., there is a single orbit, and the stabilizer of any element is trivial. For both facts, let $t_0 \in X_g$ be fixed. Since g is semisimple, so is t_0 , and this implies that the centralizer $C(t_0)$ in G is connected because G is simply-connected (a result of Steinberg; see, e.g., [127, Theorem 2.15] or [19, Theorem 3.5.6]). Since g (hence t_0) is regular, which implies in general that the connected component of $C(t_0)$ is a maximal torus (see, e.g., [12, II.12.2, Proposition]), we have $C(t_0) = T$. Also T is then the unique maximal torus containing t_0 (this is easy to see here since any such lies in $C(t_0) = T$). Now, to show transitivity of the action, assume that $t \in X_g$; then t and t_0 are conjugate, say $t = gt_0g^{-1}$. Then since $t_0 \in g^{-1}Tg$, which is a maximal torus, the observation above implies that $g^{-1}Tg = T$, hence $g \in N(T)$, showing that the image of g in W sends t_0 to t .⁴ For the last point, if $w \in N(T)$ fixes t_0 , this means $w \in C(t_0) = T$, i.e., $w = 1$ in W .
- (iv) Let L/K be the splitting field of the characteristic polynomial $\det(X - g) \in K[X]$ of g , computed as stated in $GL(d)$; then any element in X_g lies in $G(L)$ (or $T(L)$), and its coefficients generate L . This is clear because L is generated by the eigenvalues of X , and finding $t \in X_g$ amounts

⁴ In fact, any two elements of a maximal torus of any connected linear algebraic group which are G -conjugate are conjugate under $N(T)$; see [31, Corollary 0.12, (iv)].

to diagonalizing g (because T is a subgroup of diagonal matrices), so that the non-zero coefficients of t are precisely the eigenvalues.

- (v) Now fix $t_0 \in X_g$ and let σ be any K -automorphism of the separable closure K^s of K . Since $\sigma(g) = g$ and $\sigma(t_0) \in T$ because T is defined over K , it follows that $\sigma(t_0) \in X_g$. Hence, for any such σ there exists (by (iii)) a unique $w_\sigma \in W$ such that $\sigma(t_0) = w_\sigma^{-1} \cdot t_0$. The map

$$\begin{cases} \text{Gal}(K^s/K) & \rightarrow & W \\ \sigma & \mapsto & w_\sigma \end{cases} \quad (\text{E.1})$$

is a group homomorphism because by choosing a representative \dot{w}_τ of $w_\tau \in W$ lying in $N(T)(K)$, we have

$$(\sigma\tau)(t_0) = \sigma(w_\tau^{-1} \cdot t_0) = \sigma(\dot{w}_\tau^{-1} t_0 \dot{w}_\tau) = \dot{w}_\tau^{-1} \sigma(t_0) \dot{w}_\tau = w_\tau^{-1} \cdot w_\sigma^{-1} \cdot t_0.$$

By (iv), the kernel of this homomorphism is exactly $\text{Gal}(K^s/L)$, hence it induces an injective homomorphism $\text{Gal}(L/K) \rightarrow W$, as desired. \square

The restriction to regular elements can be bypassed by specialization: working with the field K which is the function field of G , there is a ‘generic’ element $\eta \in G(K)$ (for instance, if $G = SL(n)$, then η is a matrix $(t_{i,j})$ with indeterminates $t_{i,j}$ satisfying the only relation $\det(t_{i,j}) = 1$). Clearly, η is regular and semisimple, and thus we have an injection

$$\varphi : \text{Gal}(L/K) \rightarrow W$$

where L/K is the splitting field of the characteristic polynomial P_η of η (which is in $K[T]$). Any $g \in G(\mathbf{Q})$ is a specialization of η and its characteristic polynomial is a specialization of P_η ; thus the Galois group of its splitting field is isomorphic to a subgroup of W . (Note that, in fact, φ is an isomorphism for all the groups considered.)

Note also that Corvaja [25, Corollary 1.11] has proved general results showing that the ‘generic’ Galois group (that of P_η) is always ‘attained’ by some rational element $g \in G(K)$ if $G(K)$ is Zariski-dense in G , and K is finitely generated.

Exercise E.1 Here is a slightly different argument leading to the same conclusion (see Section 1.1 of the Bourbaki Seminar talk of G. Laumon on Lusztig’s character sheaves for instance). Again let G/K be a connected reductive group defined over K , and $T \subset G$ a maximal torus defined over K .

- (1) For any regular semisimple element $g \in G(K)$, let

$$Y_g = \{h \in G/T \mid h^{-1}gh \in T\}.$$

- Show that Y_g is not empty and that W acts on Y_g by $w \cdot hT = hw \cdot T$. Show also that $\text{Gal}(K^s/K)$ acts on Y_g by $\sigma(hT) = \sigma(h)T$.
- (2) Show that the action of W on Y_g is free and transitive. (This does not require G to be simply-connected.)
 - (3) Now fix $h_0 \in Y_g$ and assume G is a subgroup of $GL(r)$ for some r with T a subgroup of the torus of diagonal matrices. Let $t_0 = h_0^{-1}gh_0 \in T$. Show that the splitting field L of the characteristic polynomial of g (computed in $GL(r)$) is the field generated by the non-zero coefficients of t_0 .
 - (4) For $\sigma \in \text{Gal}(K^s/K)$, define $w_\sigma \in W$ to be the unique element such that $\sigma(h_0T) = h_0w_\sigma T$. Prove that if $C(g)$ is connected (e.g., if G is simply-connected), then $\sigma \in \text{Gal}(K^s/K)$ fixes L if and only if $w_\sigma = 1$. [Hint: Show that $\sigma(t_0) = w_\sigma^{-1}t_0w_\sigma$.]
 - (5) Deduce that under the assumption of (4), there exists an injective homomorphism $\text{Gal}(L/K) \rightarrow W$.

These arguments also provide an alternative sieve path to results such as Theorem 7.12 or those concerning Chavdarov's problem in Chapter 8. Indeed, assume G is, like $GL(n)$ or $Sp(2g)$, 'defined over \mathbf{Z} ' so that, for all primes ℓ , the algebraic group G/\mathbf{F}_ℓ over the finite field \mathbf{F}_ℓ makes sense. Then we also have the finite groups $G(\mathbf{F}_\ell)$ of rational points, and moreover we have reduction maps $G(\mathbf{Q}) \rightarrow G(\mathbf{F}_\ell)$ for all ℓ . Another basic fact is that reduction leads to a canonical isomorphism

$$W(G/\mathbf{Q}) \rightarrow W(G/\mathbf{F}_\ell).$$

On the other hand, Galois theory shows that the Galois group D_ℓ of the ℓ -adic completion of \mathbf{Q} is isomorphic to a subgroup of $\text{Gal}(\bar{\mathbf{Q}}/\mathbf{Q})$, and that there is a surjection

$$D_\ell \rightarrow \text{Gal}(\bar{\mathbf{F}}_\ell/\mathbf{F}_\ell).$$

The description of the homomorphism (E.1) associated with a given $g \in G(\mathbf{Q})$ clearly shows that the diagram

$$\begin{array}{ccc} D_\ell & \longrightarrow & W(G/\mathbf{Q}) \\ \downarrow & & \downarrow \\ \text{Gal}(\bar{\mathbf{F}}_\ell/\mathbf{F}_\ell) & \longrightarrow & W(G/\mathbf{F}_\ell) \end{array}$$

commutes. Hence the action of the Frobenius in $\text{Gal}(\bar{\mathbf{F}}_\ell/\mathbf{F}_\ell)$ gives elements in W which belong to the image of $\text{Gal}(\bar{\mathbf{Q}}/\mathbf{Q}) \rightarrow W$, i.e., to the Galois group of the splitting field of g . One can translate this back into factorization patterns (indeed, this is quite clear for $GL(n)$ with the standard diagonal torus), or one could choose to sieve using sieving sets in $G(\mathbf{F}_\ell)$ made of elements which

are conjugate to matrices in $T(\mathbf{F}_\ell)$ where the Frobenius acts like some given element w in the Weyl group.

Remark E.2 This suggests strongly to look for results like Theorem 7.12 for other groups. There are a few potential subtleties. For instance, for $G = SO(2n + 1)$, $n \geq 1$, it is well known that G is semisimple, with Weyl group $W \simeq W_{2n}$, but there is a ‘functional equation’ which imposes that 1 is a root of the characteristic polynomial of any $g \in SO(2n + 1)$, so one can expect a ‘maximal’ splitting field (generated by roots of the characteristic polynomial), in the sense of one with Galois group W_{2n} , but *not* an irreducible characteristic polynomial. This type of statement is proved in F. Jouve’s Ph.D. thesis [69] (see also the results [75] of Katz, in the setting of the sieve for Frobenius).

Also the condition of simple-connectedness is not simply technical. Indeed, consider $G = PSL(2)/\mathbf{Q}(i)$ and the class of the matrix

$$g = \begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix}$$

which is a regular element of $G(\mathbf{Q}(i))$. The centralizer of g in G is the union of the matrices of the two types

$$\begin{pmatrix} a & 0 \\ 0 & a^{-1} \end{pmatrix}, \quad \begin{pmatrix} 0 & -a \\ a^{-1} & 0 \end{pmatrix}$$

so it is of dimension 1 but not connected.⁵ It is the normalizer of the diagonal maximal torus of $PSL(2)$, so that the second component represents the non-trivial element of the Weyl group. The defining field of g is $\mathbf{Q}(i)$, but to speak of characteristic polynomial we must use a faithful representation, the simplest of which is the symmetric square $PSL(2) \rightarrow GL(3)$,⁶ which maps

$$g \mapsto \begin{pmatrix} -1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -1 \end{pmatrix}$$

and there the characteristic polynomial has trivial splitting field.

A last issue would be to consider non-split groups, e.g., unit groups of quaternion algebras for which the group of \mathbf{Q} -rational points is another very interesting type of arithmetic group. But we will leave this for the future.

⁵ This is one of the simplest examples of this phenomenon for a connected group.

⁶ Which can be seen as the action of $PSL(2)$ on quadratic polynomials $aX^2 + bXY + cY^2$ induced by unimodular linear substitutions, i.e., by $SL(2)$.

Appendix F

Probability theory and random walks

This chapter is simply a review of probabilistic language, in particular with respect to random walks. The sole intent is to define all terms that appear in the text so that each statement can be understood by readers not familiar with probability theory.

F.1 Terminology

A *probability space* is a triple $(\Omega, \Sigma, \mathbf{P})$, where Ω is a set of *elementary events*, Σ is a σ -algebra on Ω and \mathbf{P} is a measure on Σ such that $\mathbf{P}(\Omega) = 1$. A measurable subset $A \in \Sigma$ is also called an *event*, and if $\mathbf{P}(A) = 1$, then A is said to be *almost sure*.

Given $(\Omega, \Sigma, \mathbf{P})$, if (Y, \mathcal{F}) is any set with a σ -algebra \mathcal{F} , a *Y -valued random variable* is a measurable map $X : (\Omega, \Sigma) \rightarrow (Y, \mathcal{F})$. If Y is not specified, it is implicitly assumed to be $(\mathbf{R}, \mathcal{B})$, i.e., the real numbers with the Borel σ -algebra.

Quite often Ω is actually implicit. What is important are the random variables defined on Ω , which are introduced by a statement such as '*Let (X_n) be a sequence of random variables such that . . .*', with the meaning that we assume that some probability space is given on which a sequence of random variables exists with the given conditions. If this is unfamiliar, it is perfectly fine to assume, in such a situation, that $(\Omega, \Sigma, \mathbf{P})$ is the interval $[0, 1]$ with the Lebesgue measure.¹

The most basic properties of random variables are related to their distribution and their independence. The *distribution*, or *law*, of a random variable

¹ Even Brownian motion, which may naturally be seen as a probability measure on the space $\Omega = C([0, +\infty[, \mathbf{R})$ of continuous functions on $[0, +\infty[$, was first defined by N. Wiener as a 'random function' $[0, 1] \rightarrow C([0, 1], \mathbf{R})$, i.e., with $\Omega = [0, 1]$.

$X : \Omega \rightarrow Y$ is the measure $X(\mathbf{P})$ on Y defined by push-forward of the probability measure \mathbf{P} . Hence knowing $X(\mathbf{P})$ is equivalent to knowing the probabilities

$$\mathbf{P}(X \in A) = \mathbf{P}(\{\omega \in \Omega \mid X(\omega) \in A\})$$

for all $A \in \mathcal{F}$. The notation on the left, where neither Ω nor its elements are explicitly mentioned, is the standard probabilistic custom, and emphasizes the viewpoint that X , instead of a function on some big unknown space, is really a variable.

If $(X_i)_{i \in I}$ is any family of random variables (which may take values in different measure spaces), then the family is *independent* if and only if, for any finite set $J \subset I$, and any measurable sets A_j for $j \in J$ (in the target space of X_j), we have

$$\mathbf{P}(X_j \in A_j \text{ for all } j \in J) = \prod_{j \in J} \mathbf{P}(X_j \in A_j).$$

Equivalently, for any finite $J \subset I$, the law of the ‘random vector’ $X = (X_j)_{j \in J}$ is the product measure of the laws of the components:

$$X(\mathbf{P}) = \bigotimes_{j \in J} X_j(\mathbf{P}). \quad (\text{F.1})$$

Using the characteristic functions $\mathbf{1}_{A_i}$, this applies also to define a family of *independent events* $(A_i)_{i \in I}$, where it boils down to the condition

$$\mathbf{P}\left(\bigcap_{j \in J} A_j\right) = \prod_{j \in J} \mathbf{P}(A_j)$$

for any finite subset $J \subset I$.

Now a basic fact of measure theory is that a statement which begins by assuming the existence of a sequence of random variables with arbitrarily prescribed distributions and independence properties is not a statement concerning the empty set: given such distributions, there always exists a probability space Ω and random variables on it with the required properties.

In particular, in Chapter 7, we will want to start with a sequence of random variables (ξ_k) taking values in a finite subset S of a (discrete) group G , and such that the (ξ_k) are independent and have the same distribution for all k . So the previous statement says that this is always possible, whichever distributions we want to select for the variables. Again, if this is unfamiliar, this may be constructed simply using $\Omega = [0, 1]$: assume for instance (the simplest, yet

most important, example) that we want ξ_k to be *uniformly distributed* on the set S , i.e., such that

$$\mathbf{P}(\xi_k = s) = \frac{1}{|S|} \quad \text{for all } s \in S.$$

Then one (artificial but) perfectly suitable model of this is to enumerate $S = \{s_0, \dots, s_{n-1}\}$ in some way, where $n = |S|$, to take $\Omega = [0, 1]$, to look at the expansion of a real number $\omega \in [0, 1]$ in n -ary digits, say

$$\omega = 0.d_1d_2\dots d_l\dots$$

with $d_i \in \{0, \dots, n-1\}$, and to define simply $\xi_k(\omega) = s_{d_i}$.

Consider now a random variable X taking value in a finite-dimensional normed \mathbf{R} -vector space V , with the Borel σ -algebra on the latter. Then it makes sense to speak of the integrability of X , and of its integral with respect to \mathbf{P} (if $\dim V > 1$, this integral is taken coordinate-by-coordinate, after choosing a basis, and is of course independent of this choice), defined if

$$\int_{\Omega} \|X(\omega)\| d\mathbf{P}(\omega) < +\infty.$$

When X is thus integrable, the integral is called the *expectation* of X , and is denoted

$$\mathbf{E}(X) = \int_{\Omega} X(\omega) d\mathbf{P}(\omega),$$

which may also be computed by the formula

$$\mathbf{E}(X) = \int_V x d\mu(x)$$

if the distribution $\mu = X(\mathbf{P})$ of X is known (again, computed coordinate-wise). We will apply this, in addition to the standard case where $V = \mathbf{R}$ or \mathbf{C} , to situations where X takes values in a vector space of matrices, and in fact X will only take finitely many values so that integrability will not be an issue then.

In addition, if X is square-integrable (hence integrable because Ω has finite measure), its *variance* is

$$\mathbf{V}(X) = \mathbf{E}\left((X - \mathbf{E}(X))^2\right) = \mathbf{E}(X^2) - \mathbf{E}(X)^2.$$

If (X, Y) are independent V -valued random variables, with distributions $\mu = X(\mathbf{P})$ and $\nu = Y(\mathbf{P})$, then from (F.1), it follows that for *any* measurable function $g : V \times V \rightarrow \mathbf{R}$, we have

$$\mathbf{E}(g(X, Y)) = \int_V \int_V g(x, y) d\mu(x) dv(y),$$

whenever one side is well-defined (in which case the other is also).

In particular, if X and Y are independent integrable real-valued random variables, taking $g(x, y) = xy$ we find that the right-hand side splits as a product, giving

$$\mathbf{E}(XY) = \mathbf{E}(X)\mathbf{E}(Y).$$

If (X, Y) are integrable random variables with values in the space $M(n, \mathbf{R})$ of real matrices (or in $M(n, \mathbf{C})$), we find that the (i, j) -th entry of $\mathbf{E}(XY)$ is

$$\begin{aligned} \mathbf{E}(XY)_{i,j} &= \mathbf{E}((XY)_{i,j}) = \mathbf{E}\left(\sum_k X_{i,k} Y_{k,j}\right) = \sum_k \mathbf{E}(X_{i,k} Y_{k,j}) \\ &= \sum_k \mathbf{E}(X_{i,k}) \mathbf{E}(Y_{k,j}) = (\mathbf{E}(X)\mathbf{E}(Y))_{i,j}, \end{aligned}$$

(using the fact that for any i, j, k and ℓ , the components $X_{i,j}$ and $Y_{k,\ell}$ are independent real-valued random variables). In other words, the relation

$$\mathbf{E}(XY) = \mathbf{E}(X)\mathbf{E}(Y) \tag{F.2}$$

still holds, where the product is the matrix product.

F.2 The Central Limit Theorem

Among all probability distributions, the most important is without doubt the normal distribution. A real-valued random variable X follows the normal distribution with expectation m and variance $\sigma^2 > 0$ if its law is the measure on \mathbf{R} given by

$$\mu_{m,\sigma} = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{t^2}{2\sigma^2}\right) dt.$$

As the terminology suggests, we have $\mathbf{E}(X) = m$ and $\mathbf{V}(X) = \sigma^2$.

The normal distribution arises in large part because of the *Central Limit Theorem*, which we state here in a weak form (much stronger versions are known).

Theorem F.1 *Let $(\Omega, \Sigma, \mathbf{P})$ be a probability space, and let (ξ_k) be a sequence of independent square-integrable real-valued random variables with identical distribution, expectation $\mathbf{E}(\xi_k) = 0$ and variance $\mathbf{E}(X_k) = \sigma^2$. Then, as $k \rightarrow +\infty$, the sequence of random variables*

$$X_k = \frac{\xi_1 + \cdots + \xi_k}{\sqrt{k}}$$

converges in distribution to a normal distribution with expectation 0 and variance σ^2 , i.e., we have

$$\mathbf{P}\left(\frac{\xi_1 + \cdots + \xi_k}{\sqrt{k}} \in [\alpha, \beta]\right) \rightarrow \frac{1}{\sigma\sqrt{2\pi}} \int_{\alpha}^{\beta} \exp\left(-\frac{t^2}{2\sigma^2}\right) dt$$

for any fixed real numbers $\alpha < \beta$.

This basic result helps to understand (intuitively and rigorously) the behaviour of random walks on \mathbf{Z}^d , as explained below.

F.3 The Borel–Cantelli lemmas

One often encounters events A of the type: ‘infinitely many among the properties P_n hold’, where the properties P_n define a sequence of events. For instance, many analytic properties related to convergence of sequences can be brought to such a shape. It is then interesting to know the probability of A . The Borel–Cantelli lemmas are very useful tools for this purpose.

Lemma F.2 *Let (A_n) be an arbitrary sequence of events, and let*

$$A = \bigcap_{N \geq 1} \bigcup_{n \geq N} A_n$$

which is the event ‘ ω belongs to infinitely many among the (A_n) ’.

(1) *If the series*

$$\sum_{n \geq 1} \mathbf{P}(A_n)$$

converges, then $\mathbf{P}(A) = 0$; in other words, almost surely, an elementary event ω is in only finitely many A_n .

(2) *If the events (A_n) are independent, and if the series*

$$\sum_{n \geq 1} \mathbf{P}(A_n)$$

diverges, then $\mathbf{P}(A) = 1$.

We will use only the simpler part (1), which is easily proved by observing that for all $N \geq 1$, we have

$$\mathbf{P}(A) \leq \mathbf{P}\left(\bigcup_{n \geq N} A_n\right) \leq \sum_{n \geq N} \mathbf{P}(A_n) \rightarrow 0$$

as the tail of a convergent series.

F.4 Random walks

We conclude with some vocabulary from random walks. Given a probability space $(\Omega, \Sigma, \mathbf{P})$ (often left unspecified, as usual), a *random walk* (X_k) on a discrete group G is for us a sequence of G -valued random variables, such that $X_{k+1} = X_k \xi_{k+1}$ for $k \geq 0$, where the sequence (ξ_k) is a sequence of independent, identically distributed, G -valued random variables. The initial distribution X_0 is often constant (equal to 1).² We will only consider random walks where the steps ξ_k take only finitely many values $s \in G$, and where the set of those $s \in G$ with $\mathbf{P}(\xi_k = s) > 0$ (i.e., the support of the law of ξ_k) is a generating set for G . If $X_0 = 1$ and the law is uniform, i.e., $\mathbf{P}(\xi_k = s)$ is constant for all those s where it is non-zero, then the random walk is a *simple* random walk.

If (X_k) is such a random walk, a basic notion is that of *recurrence* or *transience*. Assume $X_0 = 1$. The random walk is transient if and only if, almost surely, there are only finitely many $k \geq 1$ for which $X_k = 1$, and otherwise it is recurrent; in that case, a zero-one law shows that in fact, the probability of coming back to the origin infinitely often is 1 (not some number in $]0, 1[$). More generally, a subset $A \subset G$ is *transient* if almost surely there are only finitely many k for which $X_k \in A$, and *recurrent* if almost surely there are infinitely many k for which $X_k \in A$.

The most basic example of random walk is obtained by considering $G = \mathbf{Z}^d$ for some $d \geq 1$, and taking (X_k) to be the simple random walk on G , i.e., the one defined by $X_0 = 0$ and $X_{k+1} = X_k + \xi_{k+1}$ with independent steps distributed uniformly by a choice of vector e_i of the canonical basis and direction of movement along this axis:

$$\mathbf{P}(\xi_k = \pm e_i) = \frac{1}{2d} \quad \text{for all } k \geq 0.$$

A famous result of Pólya (the first result in the theory of random walks) states that this random walk is recurrent for $d = 1$ and $d = 2$, and transient for $d \geq 3$. In fact, a much deeper result of Varopoulos states that if G admits a simple recurrent random walk with steps uniformly supported on a symmetric generating set of G , then G is either finite or has a finite index subgroup isomorphic to either \mathbf{Z} or \mathbf{Z}^2 (see, e.g., [132, Theorem 3.24]).

For the simple random walk (X_k) on \mathbf{Z}^d , we have

$$\mathbf{P}(X_k = 0) \sim c_d k^{-d/2}, \quad \text{as } k \rightarrow +\infty,$$

² There are of course more general types of random walks, multiplying on the left instead of the right, and non-identically distributed steps.

for some constant $c_d > 0$ (think that, by the Central Limit Theorem, X_k is essentially within the ball of radius \sqrt{k} centred at the origin, which contains around $k^{d/2}$ lattice points, and each is covered more or less equitably), and for two independent copies (X_k) and (Y_k) we have

$$\mathbf{P}(X_k = Y_k) \sim 2^{-d/2} c_d k^{-d/2}, \quad \text{as } k \rightarrow +\infty, \quad (\text{F.3})$$

which follows simply by writing

$$\mathbf{P}(X_k = Y_k) = \mathbf{P}(Z_{2k} = 0) \quad (\text{F.4})$$

where (Z_k) is another simple random walk on \mathbf{Z}^d , obtained by replacing the increments of (X_n) by minus those of (Y_n) for $n > k$. Both asymptotic bounds follow precisely from the local Central Limit Theorem (see, e.g., the very general result in [132, III, Corollary 13.10]).

For $d = 2$, we also have the following variant: for two independent random walks (X_k) and (Y_k) on \mathbf{Z}^2 , we have

$$\mathbf{P}(nX_k = mY_k \text{ for some } (n, m) \neq (0, 0) \in \mathbf{Z}^2) \sim c'_2 k^{-1}$$

with $c'_2 > 0$, in fact $c'_2 = 1/4$. In other words, this is the probability that an observer placed at the origin will, at time k , only see one of two particles moving independently at random on \mathbf{Z}^2 , the other one being hidden from view.

The proof of this was given by D. Khoshnevisan; we give a quick sketch. Replacing X_k and Y_k by X'_k and Y'_k obtained by a $\pi/2$ rotation, we see that $(X'_k, Y'_k) = (X'_{k,1}, X'_{k,2}, Y'_{k,1}, Y'_{k,2})$ is a vector in \mathbf{Z}^4 with four independent coordinates given by simple random walks on \mathbf{Z} . The desired probability is then

$$\mathbf{P}(\det(X'_k, Y'_k) = 0) = \mathbf{P}(X'_{k,1} Y'_{k,2} = X'_{k,2} Y'_{k,1}).$$

This can be written as

$$\frac{1}{2\pi} \int_{-\pi}^{\pi} \varphi^2(t) dt = \frac{1}{\pi} \int_0^{\pi} \left(\mathbf{E}(\cos^k(tS_k)) \right)^2 dt$$

where

$$\varphi(t) = \mathbf{E}(\exp(itX'_{k,1}X'_{k,2})) = \mathbf{E}(\exp(itY'_{k,1}Y'_{k,2}))$$

is the characteristic function (Fourier transform) of the product of the coordinates of either X'_k or Y'_k , which is easily checked (using independence) to be given by

$$\varphi(t) = \mathbf{E}(\cos^k(tS_k)),$$

with (S_k) another simple random walk on \mathbf{Z} . Then we have

$$\frac{1}{2\pi} \int_{-\pi}^{\pi} \varphi(t)^2 dt = \frac{1}{2\pi k} \int_0^{2\pi k} \left(\mathbf{E} \left(\cos^k \left(\frac{t}{\sqrt{k}} \frac{S_k}{\sqrt{k}} \right) \right) \right)^2 dt,$$

and the idea is that the Central Limit Theorem allows us to work ‘as if’ S_k/\sqrt{k} was replaced with a normal variable Z with mean 0 and variance 1, and $\cos^k(Zt/\sqrt{k})$ by $\exp(-Z^2t^2/2)$, so that the probability will be asymptotic with

$$\frac{1}{2\pi k} \int_0^\infty \mathbf{E} \left(\exp \left(-\frac{Z^2 t^2}{2} \right) \right)^2 dt = \frac{1}{2\sqrt{2\pi k}} \mathbf{E} \left(\frac{1}{\sqrt{Z^2 + W^2}} \right)$$

where W is another standard normal random variable independent of Z . Computing in polar coordinates leads to

$$\mathbf{P}(\det(X'_k, Y'_k) = 0) \sim \frac{1}{4k}, \quad \text{as } k \rightarrow +\infty.$$

We see in particular that the set

$$M = \{(ma, mb, na, nb) \in \mathbf{Z}^4 \mid (a, b) \in \mathbf{Z}^2, (m, n) \in \mathbf{Z}\}$$

is *recurrent*: indeed, the set of (a, b, a, b) is recurrent, this being equivalent by (F.4) with the fact that the origin is recurrent in a simple random walk on \mathbf{Z}^2 , which is true as we have already mentioned.

Appendix G

Sums of multiplicative functions

This appendix reviews, for completeness, some basic statements about sums of values of multiplicative functions. Recall that a function f defined for integers $n \geq 1$ is multiplicative if $f(nm) = f(n)f(m)$ for all *coprime* integers n, m . Examples are n^s , for $s \in \mathbf{C}$ arbitrary, $\varphi(n)$, $\psi(n)$, $\mu(n)$, as well as ordinary products $f(n)g(n)$ and Dirichlet convolutions

$$(f \star g)(n) = \sum_{d|n} f(d)g\left(\frac{n}{d}\right),$$

when f and g are themselves multiplicative.

G.1 Some basic theorems

The most classical result is due to Wirsing. Here is one version:

Theorem G.1 *Let f be a non-negative multiplicative function such that*

$$\sum_{p^k \leq L} f(p)^k \log p = \kappa \log L + O(1) \tag{G.1}$$

for $L \geq 2$, where $\kappa \geq 0$ is a constant. Then we have

$$\sum_{n \leq L}^b f(n) = c(\log L)^\kappa + O((\log L)^{\kappa-1})$$

for $L \geq 2$ where c is the absolutely convergent Euler product given by

$$c = \frac{1}{\Gamma(\kappa)} \prod_p (1 - p^{-1})^\kappa (1 + f(p)),$$

and the implied constant depends only on f .

For the proof, see, e.g. [67, Theorem 1.1], with $f(n)$ replaced by $\mu^2(n)f(n)$; note that the second assumption (Equation (1.89)) of [67] is valid here when summing over squarefree integers because (G.1) implies (by summation by parts)

$$\sum_{p \leq L} f(p) = \kappa \log \log L + O(1)$$

and hence (remembering that $f(n) \geq 0$) we have

$$\sum_{n \leq L}^b f(n) \leq \prod_{p \leq L} (1 + f(p)) \leq \exp \left(\sum_{p \leq L} f(p) \right) \ll (\log L)^\kappa.$$

For a version of this theorem with explicit dependency on f (which is an important issue in some applications), see, e.g. [55, Lemma 5.4].

This first result is applicable essentially when $f(p)$, for p prime, is roughly equal to κ/p , for instance if

$$f(p) = \frac{\kappa}{p} \left(1 + O \left(\frac{1}{p^\delta} \right) \right)$$

for all primes and some constant $\delta > 0$ (depending only on f , as does the implied constant), in which case the hypothesis is a consequence of the Mertens formula. In sieve methods, this corresponds to small sieves of ‘dimension κ ’. From Theorem G.1, one easily gets by positivity lower bounds such as

$$\sum_{n \leq L}^b n f(n) \gg c L (\log L)^{\kappa-1},$$

for $L \geq 2$, where the implied constant depends on f ; this is the type of estimate suitable for ‘large sieve’ situations, where the density (written $n f(n)$) has positive lower bound over primes.

The next result we quote is one of Lau and Wu [88], and generalizes partly Wirsing’s result by allowing the *summation condition* to be of the form $g(m) \leq L$ for some other multiplicative function (close to m , in some sense), instead of $m \leq L$. On the other hand, the conditions on the multiplicative function are stronger, though very reasonable in applications. We state the version adapted to a large sieve context.

Theorem G.2 *Let f and g be multiplicative functions, $f(n) \geq 0$ and $g(n) > 0$ for all $n \geq 1$, such that*

$$f(p) = \kappa + O(p^{-\eta}), \quad g(p) = \alpha p + \alpha' p^{\theta'}, \text{ for } p \text{ prime,}$$

for some constants $\kappa > 0$, $\eta > 0$, and $\alpha > 0$, $\alpha' \neq 0$, $\theta' < 1$, the implied constant depending on f .

Then we have

$$\sum_{g(m) \leq L}^b f(m) = cL(\log L)^{-1+\kappa/\alpha} + O(L(\log L)^{-2+\kappa/\alpha}(\log \log L))$$

for $L \geq 2$, where

$$c = \frac{1}{\Gamma(\kappa/\alpha)} \prod_p (1 - p^{-1})^{\kappa/\alpha} (1 + f(p)g(p)^{-1}),$$

and the implied constant depends only on f and g ; the Euler product defining c converges absolutely.

Proof This is a special case, and a weakening of the conclusion, of Theorem 1 of [88], where precisely $f(n)$ should be replaced with $\mu^2(n)f(n)$ to incorporate the restriction of the sum to squarefree numbers; the additional parameters of [88] are given by $(\theta, \theta'', \psi, C_2, C_3, t(p)) = (1, 0, 2, 0, 0, 0)$; we take $J = 0$ in the statement of Theorem 1 of [88] (see Remarks (i), (ii) just following it). Note also that the first condition in (1.2) of [88], namely $|\kappa| < \eta^{-1}$, is not necessary here (its purpose is to ensure that the implied constant in (1.6) of [88] is independent of κ , which is crucial for later applications in [88], but mostly irrelevant in the current situation). Alternatively, note that η may be replaced by any smaller (positive) number, to ensure that the condition $|\kappa| < \eta^{-1}$ holds. \square

Exercise G.1 Consider the sum (2.18). Show that there exists a constant $c > 0$ such that

$$\sum_{m \leq L}^b \frac{3^{\omega(m)} \varphi(m)}{\psi(m)} \prod_{\substack{\ell|m \\ \ell \equiv 1 \pmod{4}}} \left(1 + \frac{3}{\ell}\right) \sim cL(\log L)^2$$

as $L \rightarrow +\infty$.

G.2 An example

To give an idea of some of the techniques involved in such results, we provide a fairly complete sketch of proof of one bound used in Chapter 8.

Proposition G.3 Let $A > 0$ be a real number. We have

$$\sum_{n \leq L}^b \psi(n)^A \ll L^{A+1}$$

for $L \geq 2$, where the implied constant depends only on A .

Proof Let $f : [0, +\infty[\rightarrow [0, 1]$ be a smooth function compactly supported on $[0, 2]$ such that $f(x) = 1$ for $0 \leq x \leq 1$. By positivity, we have

$$\sum_{n \leq L}^b \psi(n)^A \leq \sum_{n \geq 1}^b \psi(n)^A f\left(\frac{n}{L}\right),$$

and we will bound this ‘smoothed’ expression.

Let $D(s)$ denote the Dirichlet generating series

$$D(s) = \sum_{n \geq 1}^b \psi(n)^A n^{-s}$$

which converges absolutely, and hence defines a holomorphic function, for $\operatorname{Re}(s) > A + 1$. By a basic form of the Mellin inversion formula, we obtain after exchanging the order of summation and integration the relation

$$\sum_{n \geq 1}^b \psi(n)^A f\left(\frac{n}{L}\right) = \frac{1}{2i\pi} \int_{(A+2)} D(s) \hat{f}(s) L^s ds$$

where $\hat{f}(s)$ is the Mellin transform of f , namely

$$\hat{f}(s) = \int_0^{+\infty} f(x) x^{s-1} dx,$$

which is a holomorphic function for $\operatorname{Re}(s) > 0$ (because of the support condition on f). Here, and further below, $\int_{(c)}$ means a complex integration over a vertical line $\operatorname{Re}(s) = c$, oriented upwards.

Moreover, by the familiar duality between smoothness of a function and decay of its Fourier or Mellin transform, for any $C > 0$, we have

$$\hat{f}(s) \ll (1 + |\operatorname{Im}(s)|)^{-C} \quad (\text{G.2})$$

for all s in a fixed strip $0 < \delta < \operatorname{Re}(s) < B$, the implied constant depending on f , C , δ and B .

On the other hand, we claim that the Dirichlet series $D(s)$ may be analytically continued to a function meromorphic in $\operatorname{Re}(s) > A + \frac{1}{2}$ with a single simple pole at $s = A + 1$, and that this continuation has polynomial growth, in the sense that for any vertical strip $A + \frac{1}{2} < A + \frac{1}{2} + \delta < \operatorname{Re}(s) < B$, we have

$$D(s) \ll 1 + |\operatorname{Im}(s)| \quad (\text{G.3})$$

for $|\operatorname{Im}(s)| \geq 1$ (to avoid the pole), where the implied constant depends on f , δ and B .

Taking this for granted, we can combine the fast decay of $\hat{f}(s)$ and the moderate growth of $D(s)$ to move the line of integration to any fixed vertical

line $\operatorname{Re}(s) = A + \frac{1}{2} + \varepsilon$ with $0 < \varepsilon < \frac{1}{2}$, picking up the residue at the single pole $s = A + 1$:

$$\frac{1}{2i\pi} \int_{(A+2)} D(s) \hat{f}(s) L^s ds = \operatorname{Res}_{s=A+1} D(s) \hat{f}(s) L^s + \frac{1}{2i\pi} \int_{\left(A+\frac{1}{2}+\varepsilon\right)} D(s) \hat{f}(s) L^s ds$$

(precisely, apply first Cauchy's residue theorem for the rectangular contour with vertices at $A + 2 \pm iT$ and $A + \frac{1}{2} + \varepsilon \pm iT$ for some $T > 1$, then use (G.2) with $C = 3$ and (G.3) to show that, as $T \rightarrow +\infty$, the contribution to the contour integral of the horizontal segments from $A + 2 + iT$ to $A + \frac{1}{2} + \varepsilon + iT$ and from $A + \frac{1}{2} + \varepsilon - iT$ to $A + \varepsilon + iT$ tends to zero).

The residue is given by

$$\operatorname{Res}_{s=A+1} D(s) \hat{f}(s) L^s = c \hat{f}(1) L^{A+1} \ll L^{A+1},$$

where c is the residue of $D(s)$ at $s = A + 1$, and the second integral is easily estimated, using once more (G.2) with $C = 3$ and (G.3):

$$\frac{1}{2i\pi} \int_{\left(A+\frac{1}{2}+\varepsilon\right)} D(s) \hat{f}(s) L^s ds \ll L^{A+\frac{1}{2}+\varepsilon} \int_{\mathbf{R}} (1 + |t|)^{-2} dt \ll L^{A+\frac{1}{2}+\varepsilon}.$$

Since ε was chosen $< \frac{1}{2}$, this is of smaller order of magnitude than the main term, and proves the Proposition.

Hence it remains to prove the claim. To do this, we start from the Euler product expansion

$$D(s) = \prod_p (1 + \psi(p)^A p^{-s}) = \prod_p (1 + (p+1)^A p^{-s}),$$

obtained from the multiplicativity of $\psi(n)$, which is valid in the region $\operatorname{Re}(s) > A + 1$ of absolute convergence. We rewrite this as follows:

$$\begin{aligned} D(s) &= \prod_p (1 + p^{A-s}) \prod_p \left(1 + \frac{1}{1 + p^{s-A}} \left(\left(1 + \frac{1}{p} \right)^A - 1 \right) \right) \\ &= \frac{\zeta(s-A)}{\zeta(2(s-A))} \prod_p \left(1 + \frac{1}{1 + p^{s-A}} \left(\left(1 + \frac{1}{p} \right)^A - 1 \right) \right) \end{aligned}$$

where $\zeta(s)$ is the Riemann zeta function, which is well known to have meromorphic continuation to \mathbf{C} with a single simple pole at $s = 1$ (with residue 1),

while $\zeta(2s)^{-1}$ is holomorphic for $\operatorname{Re}(s) > \frac{1}{2}$. Further, notice that by the mean-value theorem, we have

$$\left| \left(1 + \frac{1}{p} \right)^A - 1 \right| \leq A 2^{A-1} p^{-1}.$$

Therefore

$$1 + \frac{1}{1 + p^{s-A}} \left(\left(1 + \frac{1}{p} \right)^A - 1 \right) = 1 + O(p^{A-1-\operatorname{Re}(s)}),$$

which, by comparison with the infinite series, shows that the second term in the expression for $D(s)$, say $E(s)$, converges absolutely (hence is holomorphic) for $\operatorname{Re}(s) > A$.

To deduce (G.3), we use the fact that a Dirichlet series (here, $E(s)$) which is absolutely convergent for $\operatorname{Re}(s) > \sigma_0$ is uniformly bounded for $\operatorname{Re}(s) > \sigma_0 + \delta$ for any $\delta > 0$ (the upper bound, which is simply the value at $\sigma_0 + \delta$ of the Dirichlet series formed with absolute values of the coefficients, depending on δ). This means that it suffices to prove (G.3) for $\zeta(s-A)$ (since $\zeta(2(s-A))^{-1}$ is itself represented by an absolutely convergent Dirichlet series for $\operatorname{Re}(s) > A + \frac{1}{2}$), and this is an immediate consequence of any standard bound for $\zeta(s)$ in the critical strip; for instance,

$$\zeta(s) \ll \frac{1}{|s-1|} + \operatorname{Re}(s)^{-1} |s|$$

for $\operatorname{Re}(s) > 0$ (which is far from the truth) suffices amply, and this in turn can be proved from the integral expression

$$\zeta(s) = \frac{s}{s-1} + s \int_1^{+\infty} \{x\} x^{-s-1} ds,$$

valid for $\operatorname{Re}(s) > 0$.

□

Appendix H

Topology

This chapter is a short survey of the two basic topological invariants which are involved in some of the geometric applications of the large sieve in Chapter 7, namely the fundamental group and the first homology group of a topological space X . We also give the definition and the simplest results concerning the mapping class group of surfaces.

Throughout this appendix, X denotes a separated (Hausdorff) topological space, which has reasonable local connectivity properties; one may without loss ask that every point in X has a neighbourhood U which is homeomorphic with an open ball in some Euclidean space \mathbf{R}^d for some $d \geq 1$. (That is, X is a topological manifold without boundary; of course, much weaker assumptions are enough.)

For references in the first sections, we will use [49] and [43].

H.1 The fundamental group

The first invariant is the *fundamental group* $\pi_1(X, x_0)$ of X , relative to a base point $x_0 \in X$. As a set, this is defined as the set of *homotopy classes* of *loops* based at x_0 , i.e., the set of continuous maps

$$[0, 1] \xrightarrow{\gamma} X$$

such that $\gamma(0) = \gamma(1) = x_0$, modulo the equivalence relation (called homotopy) for which $\gamma_1 \sim \gamma_2$ if and only if there exists a continuous map

$$\gamma : [0, 1] \times [0, 1] \rightarrow X$$

with $\gamma(0, t) = \gamma_1(t)$ and $\gamma(1, t) = \gamma_2(t)$ for $0 \leq t \leq 1$, and $\gamma(u, 0) = \gamma(u, 1) = x_0$ for all u (in other words, one can deform γ_1 to γ_2 continuously, as loops based at the point x_0).

The trivial (constant) loop $\gamma_0(t) = x_0$ for all t as identity, the reversed loop $\gamma^{-1}(t) = \gamma(1 - t)$ as inverse, and the concatenation operation $\gamma = \gamma_1\gamma_2$ of two loops

$$\gamma(t) = \begin{cases} \gamma_1(2t), & \text{if } 0 \leq t \leq 1/2 \\ \gamma_2(2(t - 1/2)), & \text{if } 1/2 \leq t \leq 1 \end{cases}$$

as product, induce on this set a group structure.

Example H.1

- (1) For any $d \geq 1$ and $x_0 \in \mathbf{R}^d$, the fundamental group of \mathbf{R}^d based at x_0 is trivial: indeed, if we put

$$h(u, t) = u\gamma(t) + (1 - u)x_0,$$

we obtain a homotopy from the constant loop $h(0, t) = x_0$ to $h(1, t) = \gamma(t)$. More generally, if X is connected and $\pi_1(X, x_0)$ is trivial, it is called *simply-connected*.

- (2) If X is pathwise connected, then it follows easily that the fundamental group does not depend on the base point up to isomorphism (use a continuous path joining two given points to move loops based at one point to loops based at the other and back). Thus we sometimes write simply $\pi_1(X)$ when only the isomorphism type is of concern. It is also clear that, if X has two or more pathwise connected components, the fundamental group based at a point x_0 is the same as that of the path-connected component of x_0 (the loops can not escape to another component).
- (3) For $X = \mathbf{S}^1$ (the unit circle in the complex plane), we have $\pi_1(X, x_0) \simeq \mathbf{Z}$ for all x_0 , where a generator is the ‘obvious’ loop obtained by following the circle once in the positive direction, in other words

$$\gamma_1(t) = e^{2i\pi t}$$

using the identification with complex numbers. For a full proof of this crucial fact, see, e.g., [49, VI.2].

- (4) Let Σ_g be an orientable compact connected surface of genus $g \geq 0$ (the boundary of a ‘doughnut with g holes’, see Figure H.1). A classical result (going back at least to Poincaré) is that $\pi_1(\Sigma_g)$ is a group generated by $2g$ elements $a_1, \dots, a_g, b_1, \dots, b_g$, subject to the single relation

$$[a_1, b_1] \cdots [a_g, b_g] = 1 \tag{H.1}$$

involving their commutators. For $g = 0$, this is the trivial group. For $g = 1$, this means there are two generators which commute, so $\pi_1(\Sigma_1) \simeq \mathbf{Z}^2$, but for any $g \geq 2$, the fundamental group is non-abelian. See, e.g., [43, 17c] for details.




Figure H.1 A compact surface

In the figure, the generators are the loops starting from the base point, connecting to one of the ‘cycles’ around the holes, and coming back. (To avoid cluttering the drawing, only the connections to the first two holes are shown.)

- (5) Continuing with this example, let H_g be the ‘filled’ doughnut with g holes, or handlebody, with boundary Σ_g (intuitively, though not with the usual meaning of the word in mathematics, the ‘interior’ or ‘inside’ of Σ_g). This is a compact connected 3-manifold with boundary Σ_g . Having filled the interior, it is clear that half of the loops in a system of generators of $\pi_1(\Sigma_g)$ become trivial in $\pi_1(H_g)$ (note that here we use the map $\pi_1(\Sigma_g, x_0) \rightarrow \pi_1(H_g, x_0)$ coming ‘functorially’ from the inclusion of the boundary in H_g , see below). The relation (H.1), after replacing (say) $b_1 = \dots = b_g = 1$ becomes tautological. Indeed, one can show that $\pi_1(H_g, x_0)$ is a free group generated by the g remaining loops in Σ_g .
- (6) Let X be a compact topological manifold without boundary, of dimension $d \geq 1$ (i.e., every point has a neighbourhood homeomorphic to an open ball in \mathbf{R}^d). Then the fundamental group of X is finitely generated.

Of course, it is immediate that the fundamental group of a space is a topological invariant of this space, if it is connected: homeomorphic spaces have isomorphic fundamental groups. In particular, this means it can be used to show that certain spaces are *not* homeomorphic: for instance, $\pi_1(M(2, \mathbf{R})) \simeq \pi_1(\mathbf{R}^4) = 1$ while¹ $\pi_1(GL(2, \mathbf{R})) \simeq \mathbf{Z}$, so the ring of matrices and the group $GL(2, \mathbf{R})$ are not homeomorphic. However, more than that, one should keep in mind the more precise property that *the fundamental group is functorial* (for

¹ This follows, for instance, from the well-known homeomorphism $\mathbf{R} \times]0, +\infty[\times \mathbf{S}^1 \rightarrow SL(2, \mathbf{R})$ given by the matrix products

$$\begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix} \begin{pmatrix} a & 0 \\ 0 & a^{-1} \end{pmatrix} \begin{pmatrix} \operatorname{Re}(z) & -\operatorname{Im}(z) \\ \operatorname{Im}(z) & \operatorname{Re}(z) \end{pmatrix}.$$

topological spaces with a base point). This means that if we have topological spaces X_1 and X_2 with base points x_1 and x_2 , and a continuous map

$$X_1 \xrightarrow{f} X_2$$

with $f(x_1) = x_2$, not necessarily injective or surjective, there is always an induced map

$$f_* : \pi_1(X_1, x_1) \rightarrow \pi_1(X_2, x_2)$$

(simply obtained by letting $f_*(\gamma) = f \circ \gamma$ for any loop based at x_1 , and checking that this is compatible with the homotopy relation), and this map is itself compatible with *composition* of continuous maps (so that $(f \circ g)_* = f_* \circ g_*$, whenever $f \circ g$ is defined), and with the identity maps (so that $(\text{Id}_X)_* = \text{Id}_{\pi_1(X, x_0)}$). From this it immediately follows that if f is a homeomorphism of topological spaces, then f_* gives an isomorphism of their fundamental groups.

Much of the importance of the fundamental group lies in its relations with *coverings* of X . This is indeed where the analogy with the Galois group in algebra can be seen, and this for instance ‘explains’ why there are algebraic fundamental groups (used in Chapter 8) for algebraic varieties which, as topological spaces, do not have the connectedness properties required for the path-based definition. A *covering*

$$Y \xrightarrow{f} X$$

of X with fiber F (any set) is a topological space Y with a continuous map to X , which has the following ‘local’ structure: for any $x \in X$, there exists a neighbourhood U of x and a commutative diagram

$$\begin{array}{ccc} U \times F & \xrightarrow{j} & Y \\ p \downarrow & & \downarrow f \\ U & \xrightarrow{i} & X \end{array}$$

where i is the inclusion of U in X , $p(x, y) = x$ for $(x, y) \in U \times F$ and j is a *homeomorphism*, where F is considered as a discrete topological space. In other words, over a small enough neighbourhood of x , Y is a ‘stack’ of copies of U , indexed by the set F (which does not depend on x). In particular, f is a local homeomorphism, and Y has the same good local connectivity properties as X .

As in the case of representations of groups, it is important to define the morphisms (and isomorphisms, or automorphisms) between two coverings

$$Y_1 \xrightarrow{f_2} X$$

and

$$Y_2 \xrightarrow{f_2} X :$$

a continuous map

$$Y_1 \xrightarrow{g} Y_2$$

is a morphism of coverings if the obvious triangular diagram commutes, i.e., $f_2(g(y)) = f_1(y)$ for all $y \in Y_1$ (which means that whenever two elements of Y_1 are ‘in the same fiber’, then so also are their images).

The simplest type of covering arises when the local description actually holds globally, i.e., when Y is homeomorphic (isomorphic, really) with $X \times F$ with f given by the projection map. If this is so, the covering is called *trivial*.

A simple example of a non-trivial covering is the following: let $Y = X = \mathbf{S}^1 \subset \mathbf{C}$ and $f(z) = z^2$; here the fiber F has two elements. The fact that f is indeed a covering follows easily from the fact that a square-root function on non-zero complex numbers can be defined locally, and the fact that it is not trivial reflects the fact that it can not be defined globally.

More generally, for $n \geq 1$, defining $f_n(z) = z^n$ gives a covering of the circle. In fact, any *connected* covering of the circle is of this type, up to isomorphism.

This last assertion may bring to mind the fact that the fundamental group of \mathbf{S}^1 is \mathbf{Z} . Indeed, the two properties are related, and knowing the coverings of a (reasonable) space X is equivalent with knowing its fundamental group. Precisely, assuming (as we will do from now on) that X is connected, locally pathwise connected, and moreover that each point has a simply-connected neighbourhood,² there is a precise correspondence between (1) actions of $\pi_1(X, x_0)$ on a discrete topological space F (up to isomorphisms of group actions), and (2) coverings

$$Y \xrightarrow{f} X$$

with fiber F (up to isomorphism of coverings). Moreover, a covering is connected if and only if the associated action is transitive.³

For details, see, e.g., [49, IX.6] or [43, 13, 14], but here is an intuitive indication of how to construct the action of the fundamental group on F from a covering

$$Y \xrightarrow{f} X.$$

First, we can identify F with the fiber $f^{-1}(x_0)$, and $\pi_1(X, x_0)$ will act on this fiber, in the following way: given a loop $\gamma : [0, 1] \rightarrow X$, and an element

² For instance, this is true if every point has a neighbourhood homeomorphic to an open ball in \mathbf{R}^d for some $d \geq 1$.

³ The most intrinsic way of phrasing this is to speak of equivalence of categories, to emphasize that in addition *morphisms* between the two kinds of objects also correspond.

$y \in F = f^{-1}(x_0)$, we can use the fact that f is a local homeomorphism to ‘reproduce’ γ , or at least its ‘beginning’ in Y , starting from y . Then (by a process similar to the process of analytic continuation of analytic functions), we can follow in Y , little by little, a path ‘above’ the loop γ ; at the end of the process, the loop has come back to x_0 in X , but in Y , it did not necessarily come back exactly at the starting point x . However, the end point must remain in the same fiber, i.e., it is an element y' of $f^{-1}(x_0)$, and we put $\gamma \cdot y = y'$.

Formally, quite a bit of checking must be done to make sure that following the loop is possible (this means any $\gamma : [0, 1] \rightarrow X$ can be ‘lifted’ to $\tilde{\gamma} : [0, 1] \rightarrow Y$ such that $p \circ \tilde{\gamma} = \gamma$), and then that $\gamma \cdot y$ depends only on the class of γ in the fundamental group, and satisfies the properties of an action, etc. (see, e.g., [49, IX.1.2].)

The construction in the opposite direction (from action to covering) may be summarized briefly as follows: first, it suffices to consider a transitive action (by splitting F into orbits); then, one shows how to construct a special covering \tilde{X} corresponding to the tautological left action of $\pi_1(X, x_0)$ on itself. The point is that since any transitive action is a quotient of this tautological action, ‘functoriality’ implies that \tilde{X} will similarly suffice to construct all (connected) coverings as quotients of \tilde{X} by the action of subgroups of $\pi_1(X, x_0)$. Unsurprisingly, \tilde{X} (which is defined only up to homeomorphism) is called a *universal cover* of X , and is characterized by the property of being a connected, simply-connected covering of X (i.e., $\pi_1(\tilde{X})$ is trivial). The fundamental group $\pi_1(X, x_0)$ acts on \tilde{X} (by the same process of lifting loops to \tilde{X}), and in fact acts as automorphisms of the covering. For any connected covering

$$Y \xrightarrow{f} X,$$

there exists a unique subgroup $H \subset \pi_1(X, x_0)$, up to conjugacy, such that Y is isomorphic to \tilde{X}/H as covering of X . (See, e.g., [49, IX.5] or [43, 13c] for the construction of the universal cover; the idea is quite simple: \tilde{X} is defined as the set of homotopy classes of continuous maps $\gamma : [0, 1] \rightarrow X$ originating from x_0 in X but not necessarily looping back to x_0 , and the covering map sends γ to $\gamma(1) \in X$).

Example H.2

- (1) The map $e : \mathbf{R} \rightarrow \mathbf{S}^1$ defined by $e(t) = e^{2i\pi t}$ ‘is’ the universal cover of \mathbf{S}^1 ; indeed, it is a connected and simply-connected covering. Note that $\mathbf{S}^1 \simeq \mathbf{R}/\mathbf{Z}$, corresponding to the fact that \mathbf{Z} is the fundamental group of the circle.
- (2) Let $g \geq 0$ be an integer and let Σ_g be an orientable compact connected surface of genus g ; if $g = 0$, Σ_0 is simply-connected (it is a sphere), if

$g = 1$, Σ_1 is a torus, homeomorphic to $\mathbf{S}^1 \times \mathbf{S}^1 \simeq \mathbf{R}/\mathbf{Z} \times \mathbf{R}/\mathbf{Z}$, and the universal cover can then be described by

$$\begin{cases} \mathbf{R}^2 & \rightarrow \mathbf{R}/\mathbf{Z} \times \mathbf{R}/\mathbf{Z} \\ (x, y) & \mapsto (e(x), e(y)). \end{cases}$$

On the other hand, for $g \geq 2$, one shows that the universal cover of Σ_g is the Poincaré upper half-plane $\mathbf{H} = \{z \in \mathbf{C} \mid \text{Im}(z) > 0\}$. However, describing the universal covering map is not as easy as above!

Among all coverings, a particular type is given by *Galois coverings*, of which the universal covering is one example. By definition,

$$Y \xrightarrow{f} X$$

is a Galois covering with Galois group G if it is a covering such that G acts on Y by homeomorphisms *of the covering* (i.e., such that $f(g \cdot y) = f(y)$ for all $g \in G$ and $y \in Y$), and with $Y/G \simeq X$. This means in particular that the fiber F can be identified with G .

The construction of Galois coverings from actions of the fundamental group is particularly transparent: starting from a homomorphism $\pi_1(X, x_0) \xrightarrow{\varphi} G$, where G is any group, one can consider the covering $\tilde{X}/\text{Ker } \varphi \rightarrow X$, where \tilde{X} is the universal covering of X . This is a connected covering if and only if φ is surjective. In fact, all Galois coverings arise in this manner, and are isomorphic if and only if the two homomorphisms have conjugate kernels.

Note then the analogy with classical Galois theory: if K is a field and \overline{K} is a separable closure of K , one constructs Galois subextensions of K in \overline{K} from group homomorphisms

$$\text{Gal}(\overline{K}/K) \xrightarrow{\varphi} G$$

by considering the fixed field

$$\overline{K}^{\text{Ker } \varphi},$$

which is Galois over K with Galois group isomorphic to the image of φ (in particular, it is equal to G if φ is surjective).

This type of analogy is the basis for the theory of the algebraic fundamental groups of algebraic varieties, which is defined (as the Galois group of a field is) by means of automorphisms of suitable coverings, rather than using loops which do not make sense in the desired generality.

H.2 Homology

Homology is the second invariant we will describe. Again let X be a topological space, and let A be a ring (where $A = \mathbf{Z}$ is the most important case). The *first homology group of X with coefficients in A* , denoted $H_1(X, A)$ is defined as the quotient

$$H_1(X, A) = Z_1(X, A)/B_1(X, A)$$

where:

- The module of 1-cycles $Z_1(X, A)$ is the submodule of the free A -module generated by paths $\gamma : [0, 1] \rightarrow X$ defined by the condition that the boundary vanishes, where the boundary of a path is the formal combination $\gamma(1) - \gamma(0)$ in the free module generated by points of X (examples of elements of $Z_1(X, A)$ are *loops* with $\gamma(0) = \gamma(1)$).
- The module of 1-boundaries $B_1(X, A)$ is the A -submodule of $Z_1(X, A)$ generated by the boundaries $\partial\delta$ of ‘triangles’ in X (i.e., let $\delta : \Delta \rightarrow X$ be a continuous map from the standard triangle

$$\Delta = \{(x, y) \in \mathbf{R}^2 \mid x \geq 0, y \geq 0, x + y \leq 1\},$$

and define $\partial\delta$ as the sum in the free module generated by paths of the three sides of the triangle obtained by parametrizing each side by $[0, 1]$).

As for the case of the fundamental group, it is obvious that this defines topological invariants of X , and more precisely again, that it is *functorial* with respect to continuous maps: if we have a map

$$X \xrightarrow{f} Y,$$

we obtain an induced homomorphism

$$H_1(X, A) \xrightarrow{f_*} H_1(Y, A).$$

It is almost obvious that a loop based at some point x_0 which is homotopically trivial is also homologous to zero, and this translates to the existence of a group homomorphism

$$\pi_1(X, x_0) \rightarrow H_1(X, \mathbf{Z}).$$

In fact, Hurewicz proved that this map induces an isomorphism

$$H_1(X, \mathbf{Z}) \simeq \pi_1(X, x_0)/[\pi_1(X, x_0), \pi_1(X, x_0)]$$

(i.e., the first homology group of X is the abelianization of the fundamental group); see, e.g., [43, 12c]. Because of this and the theory of coverings, we




Figure H.2 A homology basis on a surface

see that $H_1(X, \mathbf{Z})$ ‘classifies’ those Galois coverings of X which have abelian Galois group.

In particular, $H_1(\mathbf{R}^d, \mathbf{Z}) = 0$ for $d \geq 1$, $H_1(\mathbf{S}^1, \mathbf{Z}) \simeq \mathbf{Z}$, and if Σ_g is an orientable compact connected surface of genus $g \geq 0$, we have

$$H_1(\Sigma_g, \mathbf{Z}) \simeq \mathbf{Z}^{2g}$$

(the relation $\prod [a_i, b_i] = 1$ lying already in the commutator subgroup). In Figure H.2 (compare with Figure H.1), the cycles, now unattached to a base point, form a family of generators of $H_1(\Sigma_g, \mathbf{Z})$.

Why is the case $A = \mathbf{Z}$ the most important? In fact, it happens that $H_1(X, A)$ can always be described purely algebraically from $H_1(X, \mathbf{Z})$: the *universal coefficient theorem* for the first homology group states that for any ring A , we have

$$H_1(X, A) \simeq H_1(X, \mathbf{Z}) \otimes A$$

as A -modules (see, e.g., [100, Section 55]).

In particular, $H_1(X, \mathbf{Q})$ is a \mathbf{Q} -vector space of dimension equal to the rank of $H_1(X, \mathbf{Z})$, with equality if and only if the latter is a free abelian group. This dimension is called the first Betti number of X . If X is a compact topological manifold, this dimension is finite since $\pi_1(X, x_0)$ is then a finitely generated group.

Similarly, for any prime number ℓ , we have

$$H_1(X, \mathbf{F}_\ell) \simeq H_1(X, \mathbf{Z}) \otimes \mathbf{F}_\ell \simeq H_1(X, \mathbf{Z}) / \ell H_1(X, \mathbf{Z}),$$

and this is a \mathbf{F}_ℓ -vector space of dimension equal to the sum of the rank of $H_1(X, \mathbf{Z})$ and the rank of the ℓ -primary part of the torsion subgroup of $H_1(X, \mathbf{Z})$.

H.3 The mapping class group of surfaces

Our last topic concerns the *mapping class groups* of compact surfaces. This is rather more specialized, and more subtle, than the fairly general considerations of the previous sections. In particular, the author’s knowledge is quite limited,

and we will simply give the basic definitions and state a few facts to orient the reader. The survey [65] contains many more details and is quite readable even for non-specialists, while the work-in-progress [39] is already full of enlightening information. Also, the book [38], edited by B. Farb, is a rich source of information for those readers interested in going further (or simply willing to learn some of the mathematical ideas surrounding this object).

Let Σ_g be, again, an orientable connected compact surface of genus $g \geq 1$, without boundary, assumed this time to be endowed with a smooth structure (so it makes sense to speak of differentiable maps on Σ_g , etc.). This surface is unique, up to diffeomorphism.

By definition, the *mapping class group* of Σ_g , denoted Γ_g , is the group

$$\Gamma_g = \text{Diff}^+(\Sigma_g) / \sim$$

of *isotopy classes* of orientation-preserving diffeomorphisms $\Sigma_g \rightarrow \Sigma_g$, where the isotopy relation is defined by

$$\varphi_0 \sim \varphi_1$$

if and only if there exists a smooth map

$$\psi : [0, 1] \times \Sigma_g \rightarrow \Sigma_g$$

such that (1) $\psi(0, x) = \varphi_0(x)$, and $\psi(1, x) = \varphi_1(x)$ for all $x \in \Sigma_g$; (2) for any $t \in [0, 1]$, the smooth map

$$\varphi_t : \begin{cases} \Sigma_g & \rightarrow & \Sigma_g \\ x & \mapsto & \psi(t, x) \end{cases}$$

is an orientation-preserving diffeomorphism.⁴ The group structure is induced by the composition of diffeomorphisms (so the identity and inverse are ‘the same’ as for diffeomorphisms).

This is a nice definition, but it is certainly not particularly enlightening at first. The difficulty is not illusory: only the case $g = 1$ is readily understood (see below). Still, a first grasp of the nature of this group can be derived by recalling the functoriality of the fundamental group $\pi_1(\Sigma_g, x_0)$ and of the first homology group $H_1(\Sigma_g, \mathbf{Z})$: this means in particular that diffeomorphisms of Σ_g act by automorphisms of either of these groups, and after some checking, it turns out that this provides actions of Γ_g on $H_1(\Sigma_g, \mathbf{Z})$, while the action on $\pi_1(\Sigma_g)$ is only defined up to conjugation, which means algebraically that we have a map

$$\Gamma_g \rightarrow \text{Out}(\pi_1(\Sigma_g))$$

⁴ In other words, one can say that φ_0 and φ_1 are homotopic in the space $\text{Diff}^+(\Sigma_g)$.

where the group $\text{Out}(\Sigma_g)$ is the quotient of $\text{Aut}(\pi_1(\Sigma_g))$ modulo inner automorphisms. The Dehn–Nielsen–Baer theorem (see, e.g., [39, Section 3.2]) states that if $g \geq 1$, this map is injective and its image is of index 2 in $\text{Out}(\pi_1(\Sigma_g))$ (extending Γ_g with a representative of orientation-reversing diffeomorphisms, one gets an extended mapping class group Γ_g^\pm and this is in fact isomorphic to $\text{Out}(\pi_1(\Sigma_g))$). Note that this leads to a purely group-theoretic definition of Γ_g , though not necessarily an easy one to use.

We will say more about the action on homology (which is, at root, nothing more nor less than taking a loop in Σ_g and looking at what its image under a diffeomorphism looks like . . .), which features in the applications of the large sieve in Section 7.6. This action leads to a group homomorphism

$$\rho_g : \Gamma_g \rightarrow SL(2g, \mathbf{Z})$$

(where the determinant is 1 because mapping classes preserve orientation), and this map ρ_g already provides quite a bit of information on the mapping class group.

For the ‘easy’ case $g = 1$, the situation is particularly simple because ρ_1 is an isomorphism. Note that the surjectivity is clear (because $\Sigma_2 \simeq \mathbf{R}^2/\mathbf{Z}^2$ and for any $m \in SL(2, \mathbf{Z})$, we can also see m as a linear diffeomorphism of Σ_2 , which clearly maps to itself under ρ_1); injectivity is not obvious, but note that this is also a special case of the Dehn–Nielsen–Baer theorem.

In general, ρ_g is neither injective nor surjective. The image, however, is always known: it is part of the basic theory of surfaces that there exists on $H_1(\Sigma, \mathbf{Z}) \simeq \mathbf{Z}^{2g}$ a non-degenerate alternating form (the *intersection pairing*, see, e.g., [43, 18c] for a construction)

$$H_1(\Sigma, \mathbf{Z}) \times H_1(\Sigma, \mathbf{Z}) \rightarrow \mathbf{Z}$$

which is preserved by the action of diffeomorphisms, so that the image of ρ_g necessarily lands in the symplectic group $Sp(\langle \cdot, \cdot \rangle) \simeq Sp(2g, \mathbf{Z})$ for this pairing. In the figure above, where we indicated the $2g$ standard cycles which form a basis of $H_1(\Sigma_g, \mathbf{Z})$, the pairing is uniquely determined by

$$\langle a_i, a_j \rangle = \langle b_i, b_j \rangle = 0, \quad \langle a_i, b_j \rangle = \delta(i, j)$$

(in other words, those cycles form a symplectic basis).

One then shows that

$$\Gamma_g \rightarrow Sp(2g, \mathbf{Z})$$

is onto for any $g \geq 1$. This, at least, shows that Γ_g is quite a large and complicated group, and it is all the more so because the kernel T_g of ρ_g which completes the exact sequence

$$1 \rightarrow T_g \rightarrow \Gamma_g \rightarrow Sp(2g, \mathbf{Z}) \rightarrow 1,$$

and which is called the *Torelli group*, is even more mysterious. (Only in 1983 did Johnson prove that T_g is finitely generated for $g \geq 3$; for $g = 2$, T_2 is *not* finitely generated, and for no $g \geq 3$ is it known if T_g is finitely *presented* . . .)

We now switch from this very global perspective to continue with a description of some basic elements of Γ_g (which are instrumental in actually proving much of what was stated before). This not only gives an idea of the geometric structure underlying the action of mapping classes, but it may also be used to describe a finite set of generators of Γ_g , some elements of T_g , and indeed can explain why the homology action is surjective.

The basic building blocks are the *Dehn twists* (see [65, Section 4], [39, Section 2.2]). To start, consider the annulus $A \subset \mathbf{C}$ given by $1/2 \leq |z| \leq 3/2$, and the diffeomorphism

$$A \xrightarrow{T} A$$

given by

$$z = re(\theta) \mapsto re(\theta + 2\pi(r - \tfrac{1}{2}));$$

note that the action of this map can be described informally as ‘turning’ the ‘outer’ knob $|z| = 2$ once while leaving the inner knob fixed.

One then constructs a vast quantity of mapping classes by embedding the annulus in Σ_g , and extending (smoothly) the map T thus transplanted by the identity on the rest of Σ_g . The class in Γ_g of such a map depends only on the image α in Σ_g of the unit circle $\mathbf{S}^1 \subset A$ (this is a simple closed curve in Σ_g , i.e., the image of a continuous embedding $\mathbf{S}^1 \rightarrow \Sigma_g$). Such an element is called a *Dehn twist about α* and is denoted by T_α .

It turns out that $T_\alpha \neq 1$ in Γ_g for every simple closed curve α . Moreover, one can check (at least intuitively with pictures) that $\rho_g(T_\alpha)$ satisfies

$$\rho_g(T_\alpha)([\beta]) = [\alpha] + \langle [\alpha], [\beta] \rangle [\beta]$$


for any simple closed curve β and its homology class $[\beta]$.

This formula is quite useful: first, if α is homologous to zero (i.e., it bounds a disc in Σ_g), we see that T_α acts trivially on homology, i.e., it is an element of the Torelli group (non trivial because all Dehn twists are).

Moreover, if $[\alpha]$ and $[\beta]$ are generators in a standard symplectic basis for $H_1(\Sigma_g, \mathbf{Z})$, we see that $\rho_g(T_\alpha)$ is a *symplectic transvection*.

Using specific systems of generators of the symplectic group $Sp(2g, \mathbf{Z})$, one can then construct products of Dehn twists which map to those generators and prove in this manner that ρ_g is surjective (see, e.g., [39, Section 6.1.3]).

In fact, one shows much more: Dehn twists generate Γ_g , and it suffices to use finitely many of them; for instance, the Dehn twists associated with the $3g - 1$

Figure H.3 Cycles with Dehn twists generating Γ_g

simple closed curves in Figure H.3 suffice to generate Γ_g (see [65, Theorem 4.2.D] or [39, Section 4.3.3]). Hence Γ_g is a finitely generated group; this was proved by Dehn and rediscovered independently by Lickorish. (It is in fact a finitely presented group; see the discussion in [65, Theorem 4.3.D] for this much harder fact.)

We conclude with the definition of the Thurston–Nielsen classification of diffeomorphisms of surfaces (and of mapping classes), since some of the geometric applications of the large sieve in Section 7.6 are directly concerned with this.

Let φ be an orientation-preserving diffeomorphism of Σ_g . Thurston showed that one of the following three possibilities holds:

- Some power φ^n is isotopic to the identity (in other words, φ is of finite order in Γ_g ; in fact, a theorem of Nielsen shows that φ is isotopic to a diffeomorphism which is of finite order as a diffeomorphism, see [65, Theorem 7.1.A]).
- There is a finite disjoint collection of circles C_i in Σ_g and a diffeomorphism φ' in the same mapping class as φ such that $\varphi'(C) = C$, where C is the union of the circles C_i ; such an element is called *reducible* because φ' induces diffeomorphisms of the components of $\Sigma_g - C$, which are ‘simpler’ surfaces.
- Otherwise, φ is a *pseudo-Anosov* element, and there exists a representative φ' of the mapping class φ with the following dynamical property, which we describe informally since the precise definition is rather involved (see [20, Section 6]): for all but finitely many singularities (of a very special type), through each point $x \in \Sigma_g$, there pass two (smooth) curves L^+ , L^- , the leaves of the so-called expanding and contracting foliations associated with φ' , and φ' acts by ‘dilation’ with some factor $\lambda > 1$ in the direction of the expanding foliation, and by ‘contraction’ with factor λ^{-1} in the direction of the contracting foliation. The *dilation factor* is known to be an algebraic integer (in fact, a unit of a totally real number field of degree at most $2g$).

(One can also easily define pseudo-Anosov elements as those which are neither of finite order, nor reducible; then one *misses* their structural properties,

but this is of course a potentially easy way to construct them, and indeed this is how we obtain them in great abundance in Section 7.6.)

There is an alternative characterization, which is not suitable as a definition (depending as it does on some further difficult results of Thurston), but may carry more familiarity for some readers: assume $g \geq 2$, let φ be a diffeomorphism of Σ_g , and let M_φ be the *mapping torus*⁵ of φ defined by

$$M_\varphi = (\Sigma_g \times [0, 1]) / \sim$$

where the equivalence relation \sim identifies $(x, 0)$ with $(\varphi(x), 1)$.

Note that M_φ is an orientable compact connected 3-manifold, of a rather special type: it is equipped with a smooth map

$$M_\varphi \xrightarrow{f} \mathbf{S}^1$$

in such a way that all fibers $f^{-1}(t)$ are naturally diffeomorphic to Σ_g . Moreover, M_φ depends only on the mapping class of φ in Γ_g up to diffeomorphism. Then φ is:

- Of finite order if and only if M_φ has a Riemannian metric g such that M_φ is locally isometric with the 3-manifold $\mathbf{H} \times \mathbf{R}$ (where \mathbf{H} is the hyperbolic plane).
- Reducible if and only if there exists an embedding $(\mathbf{R}/\mathbf{Z})^2 \rightarrow M_\varphi$ of a torus in M_φ such that the induced map

$$\mathbf{Z}^2 \simeq \pi_1((\mathbf{R}/\mathbf{Z})^2) \rightarrow \pi_1(M_\varphi)$$

is injective (this does not depend on the base points, of course).

- Pseudo-Anosov if and only if the manifold M_φ has a hyperbolic structure: there exists a Riemannian metric g on M_φ such that M_φ is locally isometric with the hyperbolic 3-space.

Example H.3

- (1) Examples of finite order elements are fairly easy to visualize when the surface is drawn with an obvious ‘symmetry’, or as automorphisms for a complex structure on Σ_g (since automorphism groups of compact Riemann surfaces are finite). Examples of reducible elements are Dehn twists. Note that the classes of finite order elements and reducible elements are *not* disjoint.

⁵ Although the definition involves the same ingredients as that of 3-manifolds by Heegaard splitting used in Section 7.6, this is a very different (much simpler) construction.

(2) Again, for $g = 1$ we can ‘see’ clearly this classification using the fact that $\Gamma_2 = SL(2, \mathbf{Z})$. One finds (not surprisingly!) that a diffeomorphism φ of Σ_2 (for instance, an element of $SL(2, \mathbf{Z})$) is:

- *of finite order* if and only if $\rho_2(\varphi)$ is *elliptic* in $SL(2, \mathbf{Z})$, i.e., when acting on \mathbf{C} , it has two complex conjugate fixed points; an example is the matrix

$$\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix};$$

- *reducible* if and only if $\rho_2(\varphi)$ is *parabolic*, i.e., it has a single fixed point in \mathbf{R} or at ∞ ; typical elements of this type are matrices

$$\begin{pmatrix} 1 & n \\ 0 & 1 \end{pmatrix}$$

for $n \in \mathbf{Z}$;

- *pseudo-Anosov* if and only if $\rho_2(\varphi)$ is *hyperbolic*, i.e., it has two distinct real fixed points. When this is so, the matrix is diagonalizable over \mathbf{R} and has two distinct real eigenvalues, with product equal to 1, so they can be written (λ, λ^{-1}) for a unique real number $\lambda > 1$; since it is the largest root of the characteristic polynomial of $\rho_2(\varphi)$, an integral monic polynomial of degree 2, it is a unit in a real quadratic number field.

If φ is itself linear, then the dynamical structure is easily described: the expanding (respectively contracting) foliation of φ acting on $\Sigma_2 = \mathbf{R}^2/\mathbf{Z}^2$ is the foliation by images modulo \mathbf{Z}^2 of affine lines with direction given by the λ -eigenspace of $\rho_2(\varphi)$ (respectively λ^{-1} -eigenspace).

References

- [1] J. Achter: *Divisibility of function field class numbers*, preprint (2006), arXiv: math.NT/0602114.
- [2] J. Achter: *Exceptional covers of surfaces*, to appear in Canadian Math. Bull.; see also arXiv:0707.2612.
- [3] J. Achter and R. Pries: *The integral monodromy of hyperelliptic and trielliptic curves*, Math. Annalen 338 (2007), 187–206; see also arXiv:math.AG/0608038.
- [4] J. Achter and R. Pries: *Monodromy of the p -rank strata of the moduli space of curves*, preprint (2007), arXiv:0707.2110.
- [5] J.E. Andersen: *Mapping class groups do not have Kazhdan's Property (T)*, preprint (2007), arXiv:0706.2184.
- [6] M. Ayad: *Périodicité (mod q) des suites elliptiques et points S -entiers sur les courbes elliptiques*, Ann. Inst. Fourier 43 (1993), 585–618.
- [7] H. Bass, J. Milnor and J.-P. Serre: *Solution of the congruence subgroup problem for SL_n ($n \geq 3$) and Sp_{2n} ($n \geq 2$)*, Publ. Math. IHES 33 (1967), 59–137.
- [8] H. Behr: *Eine endliche Präsentation der symplektischen Gruppe $Sp_4(\mathbf{Z})$* , Math. Z. 141 (1975), 47–56.
- [9] B. Bekka, P. de la Harpe and A. Valette: *Kazhdan's Property (T)*, New Math. Monographs 11, Cambridge University Press (to appear, 2008); draft available at www.unige.ch/math/biblio/preprint/2006/BCHnouveau.pdf.
- [10] P. Billingsley: *Prime numbers and Brownian motion*, Amer. Math. Monthly 80 (1973), 1099–1115.
- [11] E. Bombieri: *Le grand crible dans la théorie analytique des nombres*, Astérisque 18, SMF 1974.
- [12] A. Borel: *Linear algebraic groups*, 2nd edition, GTM 126, Springer 1991.
- [13] J. Bourgain and A. Gamburd: *Uniform expansion bounds for Cayley graphs of $SL_2(\mathbf{F}_p)$* , Ann of Math. 167 (2008), 625–642.
- [14] J. Bourgain, A. Gamburd and P. Sarnak: *Sieving and expanders*, C. R. Acad. Sci. Paris 343 (2006), 155–159.
- [15] J. Bourgain, A. Gamburd and P. Sarnak: *Sieving, expanders and sum-product*, preprint (2007).
- [16] R. Brooks: *On the angles between certain arithmetically defined subspaces of \mathbf{C}^n* , Annales Inst. Fourier 37 (1987), 175–185.

- [17] J. Brüdern and E. Fouvry: *Lagrange's four squares theorem with almost prime variables*, J. reine angew. Math. 454 (1994), 59–96.
- [18] M. Burger: *Kazhdan constants for $SL(3, \mathbf{Z})$* , J. reine angew. Math. 413 (1991), 36–67.
- [19] R.W. Carter: *Finite groups of Lie type*, Wiley Interscience 1985.
- [20] A. Casson and S. Bleiler: *Automorphisms of surfaces after Nielsen and Thurston*, LMS Student Texts 9, Cambridge University Press 1988.
- [21] Z. Chatzidakis, L. van den Dries and A. Macintyre: *Definable sets over finite fields*, J. reine angew. Math. 427 (1992), 107–135.
- [22] N. Chavdarov: *The generic irreducibility of the numerator of the zeta function in a family of curves with large monodromy*, Duke Math. J. 87 (1997), 151–180.
- [23] L. Clozel: *Démonstration de la conjecture (τ)* , Invent. math. 151 (2003), 297–328.
- [24] A. Cojocaru: *Reductions of an elliptic curve with almost prime orders*, Acta Arithmetica 119 (2005), 265–289.
- [25] P. Corvaja: *Rational fixed points for linear group actions*, Ann. Scuola Normale Sup. Pisa (to appear); arXiv:math/0610661v2.
- [26] H. Davenport and H. Halberstam: *The values of a trigonometrical polynomial at well-spaced points*, Mathematika 13 (1966), 91–96.
- [27] S. Davis, W. Duke and X. Sun: *Probabilistic Galois theory of reciprocal polynomials*, Exposition. Math. 16 (1998), 263–270.
- [28] P. Deligne: *La conjecture de Weil, II*, Publ. Math. IHES 52 (1981), 313–428.
- [29] P. Deligne: *Cohomologie étale*, S.G.A 4 $\frac{1}{2}$, LNM 569, Springer Verlag 1977.
- [30] P. Deligne and G. Lusztig: *Representations of reductive groups over finite fields*, Ann. of Math. 103 (1976), 103–161.
- [31] F. Digne and J. Michel: *Representations of finite groups of Lie type*, LMS Student Texts 21, Cambridge University Press 1991.
- [32] W. Duke: *Elliptic curves with no exceptional primes*, C. R. Acad. Sci. Paris 325 (1997), 813–818.
- [33] W. Duke, Z. Rudnick and P. Sarnak: *Density of integer points on affine homogeneous varieties*, Duke Math. J. 71 (1993), 143–179.
- [34] N. Dunfield and W. Thurston: *Finite covers of random 3-manifolds*, Invent. math. 166 (2006), 457–521.
- [35] N. Dunfield and W. Thurston: *The virtual Haken conjecture: Experiments and examples*, Geom. Topol. 7 (2003), 399–441.
- [36] M. Einsiedler, G. Everest and T. Ward: *Primes in elliptic divisibility sequences*, LMS J. Comput. Math. 4 (2001), 1–13.
- [37] A. Eskin, S. Mozes and N. Shah: *Unipotent flows and counting lattice points on homogeneous varieties*, Ann. of Math. 143 (1996), 253–299.
- [38] B. Farb (editor): *Problems on mapping class groups and related topics*, Proc. Symp. Pure Math. 74, AMS 2006.
- [39] B. Farb and D. Margalit: *A primer on mapping class groups*, working draft v2.95 (27 Aug., 2007), available online at www.math.utah.edu/~margalit/primer/.
- [40] A. Fathi, F. Laudenbach and V. Poénaru: *Travaux de Thurston sur les surfaces*, Astérisque 66/67, SMF 1979.

- [41] W. Feller: *An introduction to probability theory and its applications*, Vol. I, 3rd edition, John Wiley & Sons 1968.
- [42] E. Fouvry and P. Michel: *Sur les changements de signe des sommes de Kloosterman*, Ann. of Math. 165 (2007), 675–715.
- [43] W. Fulton: *Algebraic topology*, GTM 153, Springer 1995.
- [44] W. Fulton and J. Harris: *Representation theory*, GTM 129, Springer 1991.
- [45] H. Furstenberg: *Noncommuting random products*, Trans. Amer. Math. Soc. 108 (1963), 377–428.
- [46] P. X. Gallagher: *The large sieve and probabilistic Galois theory*, in Proc. Sympos. Pure Math., Vol. XXIV, Amer. Math. Soc. (1973), 91–101.
- [47] A. Gamburd: *On the spectral gap for infinite index ‘congruence’ subgroups of $SL_2(\mathbf{Z})$* , Israel J. Math. 127 (2002), 157–200.
- [48] K. Girstmair: *Linear dependence of zeros of polynomials and construction of primitive elements*, Manuscripta Math. 39 (1982), 81–97.
- [49] C. Godbillon: *Éléments de topologie algébrique*, Hermann 1971.
- [50] R. Gow: *Properties of the characters of the finite general linear group related to the transpose-inverse involution*, Proc. London Math. Soc. 47 (1983), 493–506.
- [51] A. Grothendieck: *Formule de Lefschetz et rationalité des fonctions L*, Séminaire Bourbaki 290 (1964–65) 31–45 in ‘Dix exposés sur la cohomologie des schémas’, 1–15, North Holland 1968.
- [52] U. Hadad: *Uniform Kazhdan constants for $\{SL_n(\mathbf{Z})\}_{n \geq 3}$* , J. of Algebra 318 (2007), 607–618. arXiv:math.RT/0612390.
- [53] C. Hall: *Big orthogonal or symplectic monodromy mod ℓ* , Duke Math. J. 141 (2008), 179–203. arXiv:math.NT/0608718.
- [54] A. J. Hahn and O. T. O’Meara: *The classical groups and K-theory*, Grundlehren der Math. Wiss. 291, Springer-Verlag 1989.
- [55] H. Halberstam and H. Richert: *Sieve methods*, Academic Press 1974.
- [56] G. H. Hardy and E. M. Wright: *An introduction to the theory of numbers*, 5th edition, Oxford Univ. Press 1979.
- [57] P. de la Harpe: *Topics in geometric group theory*, Chicago Lectures in Math., Univ. of Chicago Press 2000.
- [58] P. de la Harpe and A. Valette: *La propriété (T) de Kazhdan pour les groupes localement compacts*, Astérisque 175, SMF 1989.
- [59] H. Helfgott: *Growth and generation in $SL_2(\mathbf{Z}/p\mathbf{Z})$* , Ann of Math. 167 (2008), 601–623. arXiv:math/0509024.
- [60] H. Helfgott and A. Venkatesh: *How small must ill-distributed sets be?*, to appear in ‘Analytic Number Theory: essays in honor of Klaus Roth’, Cambridge University Press.
- [61] S. Hoory, N. Linial and A. Wigderson: *Expander graphs and their applications*, Bull. AMS 43 (2006), 439–561.
- [62] H. Hubrechts: *Point counting in families of hyperelliptic curves*, to appear in Foundations of Comput. Math.; see also arXiv.org:math/0601438.
- [63] M. N. Huxley: *The large sieve inequality for algebraic number fields*, Mathematika 15 (1968), 178–187.
- [64] L. Illusie: *Théorie de Brauer et caractéristique d’Euler-Poincaré*, in ‘caractéristique d’Euler-Poincaré (Séminaire ENS, 1978–1979)’, 161–172, Astérisque 82–83, SMF 1981.

- [65] N. Ivanov: *Mapping class groups*, in ‘Handbook of geometric topology’, 523–633, North Holland 2002.
- [66] H. Iwaniec: *Introduction to the spectral theory of automorphic forms*, Biblioteca de la Revista Mat. Iberoamericana 1995.
- [67] H. Iwaniec and E. Kowalski: *Analytic number theory*, Colloquium Publ. 53, AMS 2004.
- [68] N. Jones: *Almost all elliptic curves are Serre curves*, preprint (2006), arXiv: math.NT/0611096.
- [69] F. Jouve: *Sommes exponentielles, crible, et variétés sur les corps finis*, Ph.D. thesis (Université Bordeaux I, December 2007).
- [70] M. Kassabov: *Kazhdan constants for $SL_n(\mathbf{Z})$* , Int. J. Algebra Comput. 15 (2005), 971–995.
- [71] N. Katz: *Sums of Betti numbers in arbitrary characteristic*, Finite Fields Appl. 7 (2001), no. 1, 29–44.
- [72] N. Katz: *Gauss sums, Kloosterman sums and monodromy*, Annals of Math. Studies 116, Princeton Univ. Press 1988.
- [73] N. Katz: *Twisted L-functions and monodromy*, Annals of Math. Studies 150, Princeton Univ. Press 2002.
- [74] N. Katz: *Moments, monodromy and perversity: a diophantine perspective*, Annals of Math. Studies 159, Princeton Univ. Press 2005.
- [75] N. Katz: *Report on the irreducibility of L-functions*, to appear (volume in honor of S. Lang).
- [76] A. J. de Jong and N. Katz: *Monodromy and the Tate conjecture: Picard numbers and Mordell-Weil ranks in families*, Israel J. Math. 120 (2000), part A, 47–79.
- [77] N. Katz and P. Sarnak: *Random matrices, Frobenius eigenvalues and monodromy*, Colloquium Publ. 45, AMS 1999.
- [78] A. A. Klyachko: *Models for complex representations of the groups $GL(n, q)$* , Mat. Sb. (1983), 371–386.
- [79] A. Knap: *Representation theory of semisimple groups*, Princeton Math. Series 36, Princeton Univ. Press 1986.
- [80] E. Kowalski: *The large sieve, monodromy and zeta functions of curves*, J. reine angew. Math 601 (2006), 29–69.
- [81] E. Kowalski: *On the rank of quadratic twists of elliptic curves over function fields*, International J. Number Theory 2 (2006), 267–288.
- [82] E. Kowalski: *Big symplectic monodromy: a theorem of C. Hall*, note available at www.math.ethz.ch/~kowalski/notes-unpublished.html.
- [83] E. Kowalski: *Exponential sums over definable subsets of finite fields*, Israel J. Math. 160 (2007), 219–252.
- [84] E. Kowalski: *Weil numbers generated by other Weil numbers and torsion fields of abelian varieties*, J. London Math. Soc. 74 (2006), 273–288.
- [85] M. Larsen: *Maximality of Galois actions for compatible systems*, Duke Math. J. 80 (1995), 601–630.
- [86] S. Lang: *Algebra*, 2nd edition, Addison-Wesley 1984.
- [87] M. Larsen: *Navigating the Cayley graph of $SL_2(\mathbf{F}_p)$* , International Math. Res. Notices 27 (2003), 1465–1471.
- [88] Y. K. Lau and J. Wu: *Sums of some multiplicative functions over a special set of integers*, Acta Arith. 101 (2002), 365–394.

- [89] Y.-R. Liu and R. Murty: *Sieve methods in combinatorics*, J. Comb. Theory, Series A, 111 (2005) 1–23.
- [90] Q. Liu: *Algebraic geometry and arithmetic curves*, Oxford Grad. Texts in Math. 6, Oxford Univ. Press 2002.
- [91] A. Lubotzky: *Discrete groups, expanding graphs and invariant measures*, Progr. Math. 125, Birkhäuser 1994.
- [92] A. Lubotzky and D. Segal: *Subgroup growth*, Progr. Math. 212, Birkhäuser 2003.
- [93] A. Lubotzky and A. Zuk: *On Property (τ)*, to appear, see www.ma.huji.ac.il/~alexlub/BOOKS/On%20property/On%20property.pdf.
- [94] G. Lusztig: *Representations of finite Chevalley groups*, CBMS 39, AMS 1978.
- [95] R. Lyndon and P. Schupp: *Combinatorial group theory*, Ergebnisse Math. 89, Springer-Verlag 1977.
- [96] J. Maher: *Random walks on the mapping class group*, preprint (2006), [arXiv: math.GT/060443](http://arxiv.org/abs/math.GT/060443).
- [97] J. Milne: *Etale cohomology*, Princeton Math. Series 33, Princeton Univ. Press 1980.
- [98] H. L. Montgomery: *The analytic principle of the large sieve*, Bull. AMS 84 (1978), 547–567.
- [99] H. L. Montgomery and R. C. Vaughan: *Multiplicative number theory I. Classical theory*, Cambridge Studies in Advanced Math. 97, Cambridge University Press 2006.
- [100] J. Munkres: *Elements of algebraic topology*, Addison-Wesley 1984.
- [101] M. Neuhauser: *Kazhdan's Property T for the symplectic group over a ring*, Bull. Belg. Math. Soc. Simon Stevin 10, no. 4 (2003), 537–550.
- [102] M. Newman: *Integral matrices*, Pure and Applied Math. 45, Academic Press 1972.
- [103] O. T. O'Meara: *Symplectic groups*, Math. Surveys 16, AMS (1978).
- [104] R. Pink: *The Mumford-Tate conjecture for Drinfeld-modules*, Publ. Res. Inst. Math. Sci. 33 (1997), 393–425.
- [105] B. Poonen: *Bertini theorems over finite fields*, Ann. of Math. 160 (2004), 1099–1127.
- [106] O. Ramaré: *Eigenvalues in the large sieve inequality*, Funct. Approx. Comment. Math 37 (2007), 399–427.
- [107] K. Ribet: *On ℓ -adic representations attached to modular forms*, Invent. math. 28 (1975), 245–275.
- [108] I. Rivin: *Counting reducible matrices, polynomials, and surface and free group automorphisms*, Duke Math. J., to appear; see [arXiv:math.NT/0604489](http://arxiv.org/abs/math.NT/0604489).
- [109] M. Rubinstein and P. Sarnak: *Chebyshev's bias*, Experimental Math. 3 (1994), 173–197.
- [110] I. Ruzsa: *Erdős and the integers*, J. Number Theory 79 (1999), 115–163.
- [111] L. Saloff-Coste: *Random walks on finite groups*, in 'Probability on discrete structures', 263–346, Encyclopaedia Math. Sci. 110, Springer 2004.
- [112] P. Sarnak: *Some applications of modular forms*, Cambridge Tracts in Math. 99, Cambridge University Press 1990.
- [113] P. Sarnak: *Slides for Rademacher Lectures*, Philadelphia, September 2006, available at <http://www.math.princeton.edu/sarnak/rademacher1.pdf>.

- [114] P. Sarnak and X.X. Xue: *Bounds for multiplicities of automorphic representations*, Duke Math. J. 64 (1991), 207–227.
- [115] J.-P. Serre: *Linear representations of finite groups*, Grad. Texts in Math. 42, Springer Verlag 1977.
- [116] J.-P. Serre: *Zeta and L-functions*, in ‘Arithmetical algebraic geometry’, 82–92 Harper and Row 1965.
- [117] J.-P. Serre: *Spécialisation des éléments de $\mathrm{Br}_2(\mathbf{Q}(T_1, \dots, T_n))$* , C. R. Acad. Sci. Paris 311 (1990), 397–402.
- [118] Y. Shalom: *Bounded generation and Kazhdan’s property (T)*, Publ. Math IHES 90 (1999), 145–168.
- [119] Y. Shalom: *The algebraization of Kazhdan’s property (T)*, ICM 2006 Madrid Proceedings, 1283–1310.
- [120] G. Shimura: *Introduction to the arithmetic theory of automorphic functions*, Princeton Univ. Press 1971.
- [121] L. Silberman: *Addendum to ‘Random walks in random groups’ by M. Gromov*, GAFA 13 (2003), 147–177.
- [122] J. Silverman: *Wieferich’s criterion and the abc-Conjecture*, J. Number Theory 30 (1988), 226–237.
- [123] J. Silverman: *p-adic properties of division polynomials and elliptic divisibility sequences*, Math. Ann. 332 (2005), 443–471; Addendum, Math. Ann. 332 (2005), no. 2, 473–474.
- [124] J. Silverman: *The arithmetic of elliptic curves*, GTM 106, Springer 1986.
- [125] T. A. Springer: *Linear algebraic groups*, 2nd edition, Progr. Math. 9, Birkhäuser 1998.
- [126] R. Steinberg: *The representations of $GL(3, q)$, $GL(4, q)$, $PGL(3, q)$ and $PGL(4, q)$* , Canad. J. Math. 3 (1951), 225–235.
- [127] R. Steinberg: *Torsion in reductive groups*, Advances in Math. 15 (1975), 63–92.
- [128] F. Taherkhani: *The Kazhdan property of the mapping class group of closed surfaces and the first cohomology group of its cofinite subgroups*, Experiment. Math. 9 (2000), 261–274.
- [129] C. R. Vinroot: *Twisted Frobenius-Schur indicators of finite symplectic groups*, J. Algebra 293 (2005), 279–311.
- [130] M. Ward: *Memoir on elliptic divisibility sequences*, Amer. J. Math. 70 (1948), 31–74.
- [131] M. Weber: *Divisors of Bernoulli sums*, Results in Mathematics 51 (2007), 141–179. arXiv:math.PR/0703696.
- [132] W. Woess: *Random walks on infinite graphs and groups*, Cambridge Tracts in Math. 138, Cambridge University Press 2000.

Index

- (T)-constant, 119, 128, 145, 146, 150, 233, 244
- (τ)-constant, 107, 119, 128, 145, 237
- L -function, 163, 171, 183
- S -integral point, 62, 65
- S -unit, 64
- \mathbf{F}_q -rank, 76
- \mathbf{F}_ℓ -adic sheaves, 2
- ℓ -adic sheaves, 154, 160, 164, 168, 173, 175
- p -primary part, 95
- q -symplectic polynomial, 158, 160, 211, 215, 216

- abelianization, 95, 96, 98, 136, 138, 140, 144, 227, 234, 237, 275
- additive character, 12, 48, 98
- additive group, 114, 246
- algebraic curve, 155, 157, 163, 168, 170, 189, 190, 193
- algebraic family of curves, 158, 159, 175, 178, 190
- algebraic variety, 160, 168, 218, 245
- alternating pairing, 176, 246, 278
- arithmetic Frobenius automorphism, 161
- arithmetic function, 1
- arithmetic fundamental group, 160, 161
- arithmetic group, 101, 132, 151
- arithmetic monodromy group, 162
- arithmetic transition, 145, 150
- automorphic form, 235

- base point, 268
- Bernoulli distribution, 89, 93
- Betti number, 168
- bipartite graph, 112, 121
- Borel subgroup, 77, 248

- Borel–Cantelli lemma, 132, 134, 258
- bounded elementary generation, 238, 243
- Brauer group, 67
- Brownian motion, 88, 254
- Brun–Titchmarsh inequality, 25, 92
- Buchstab identity, 200

- canonical height, 60, 61, 64
- Cayley graph, 110, 111, 121–123
- Central Limit Theorem, 25, 92, 93, 257, 260
- centralizer, 212
- character, 32, 34, 40, 71, 75, 167, 223–225, 227
- character group, 80, 227
- character table, 81, 82, 227, 228, 231
- characteristic function, 22, 23
- characteristic polynomial, xvi, 53, 101, 129, 133, 135, 147, 207, 211–213
- Chebotarev density theorem, 67, 202
- Chebychev estimate, 98, 143
- Chinese Remainder Theorem, 12, 31, 57, 117
- class function, 81, 223–225, 227
- classical large sieve, 30, 48, 63, 92, 102
- classical sieve setting, 9, 101
- classical sieve theory, 1, 51, 199
- closed points, 163
- closed surface, 6, 133, 137, 269, 273, 276, 277
- closed-point sieve, 67
- cocharacter group, 75, 80
- coinvariant, 166, 171, 215
- combinatorial sieve, 52, 92, 197
- commutator relation, 114, 115, 123, 146
- commutator subgroup, 113, 121, 125, 276
- compatible system, 162–164, 169, 171, 173, 176, 179, 188, 191
- complementary series, 235, 236

- conductor, 57
- congruence quotient, 81
- congruence subgroups, 113, 116, 117, 119, 120, 132, 236, 237
- conjugacy classes, 32
- conjugacy sieve, 32, 34, 40, 67, 106, 116, 122, 125, 126
- connected component of the identity, 247
- correlation coefficient, 46
- coset sieve, 36, 40, 70, 126, 154, 161, 162, 202
- coupon collector problem, 149, 150
- covering, 271–273
- crible étrange, 1
- cut-off phenomenon, 148
- cycle type, 54, 180, 204
- cyclotomic polynomial, 134

- definable set, 218
- degree, 161, 163, 194
- Dehn twist, 279, 281
- Deligne–Lusztig character, 70, 74, 75, 80, 81, 86, 227
- density, 2, 10, 26, 28, 32, 37, 46, 69, 181, 198, 204, 210, 218
- diagonal term, 29
- dihedral group, 39
- dilation factor, 135, 280
- dimension of a representation, 220
- dimension of a sieve, 200
- dimension of an irreducible representation, 70, 85, 174
- Dirichlet character, 57
- discrete group, 101, 234, 259
- discrete series, 228, 235
- discrete subgroup, 235
- distribution of a random variable, 254
- divisor, 194
- dual group, 86
- dual sieve, 22, 68, 97, 142, 200
- dual sums, 18
- duality principle, 18

- eigenvalue of the Frobenius, 166
- elementary matrix, 6, 114, 136, 145, 238, 243
- elementary subgroup, 114
- elliptic curve, 5, 59, 64, 160, 201
- elliptic divisibility sequence, 65, 66
- elliptic sieve, 59, 98
- equidistribution, 2, 20, 27, 31, 40, 51, 90, 96, 102, 106, 119, 140, 142, 200–202
- equidistribution remainder terms, 21

- étale cohomology group, 166, 169, 195
- étale covering, 67, 161, 169
- Euler product, 262
- Euler–Poincaré characteristic, 163, 165, 170, 171, 173, 179
- Euler–Poincaré formula, 170
- exceptional eigenvalue, 81, 236
- exceptional isomorphism, 115
- expanders, 112, 119, 123, 140
- expectation, xv, 89, 98, 108, 149, 256
- exponential distribution, 144
- exponential sums, 2, 22, 34, 51, 89, 99, 106, 117, 127, 164, 201, 210
- external tensor product, 34, 72, 221

- factorization of a polynomial over a finite field, 54, 204
- finite group of Lie type, 72, 101, 119, 172, 211, 224, 227, 231
- finite simple group, 152, 188
- finitely generated group, 101, 103, 105, 106, 111, 151, 233, 234, 236, 270, 276, 280
- finitely presented group, 94, 140, 280
- first Betti number, 137–140, 276
- first homology group, 133, 138, 143, 144, 268, 275, 277
- Fourier coefficients of automorphic forms, 3
- free group, 94, 136, 234, 270
- Frobenius automorphism, 54, 80, 194
- Frobenius conjugacy class, 2, 154, 163
- Frobenius reciprocity, 36, 73, 222, 226
- functoriality, 270, 273, 275, 277
- fundamental group, 94, 137, 138, 140, 144, 161, 268, 270, 272, 273, 275, 277
- fundamental lemma, 199

- Galois covering, 161, 162, 169, 274
- Galois group, 52, 54, 154, 157–159, 161, 162, 169, 170, 177, 179, 186, 190, 194, 271
- general position, 77, 80, 81
- generalized character, 75
- Generalized Riemann Hypothesis, 25, 67, 154, 201
- genus, 156, 159, 163, 175, 178, 194, 269, 276, 277
- geometric conjugacy, 76, 81, 85
- geometric conjugacy class, 77, 78, 86
- geometric distribution, 149
- geometric Frobenius automorphism, 67, 161
- geometric fundamental group, 160, 167, 169, 177, 178
- geometric generic point, 160, 189

- geometric monodromy group, 162, 177, 178, 183, 187, 189, 191
- geometrically connected, 246
- geometrically irreducible, 155, 169
- geometrically trivial covering, 161
- girth, 122
- global Frobenius automorphism, 165
- Goursat–Ribet lemma, 153, 177
- graph, 110
- Grothendieck–Lefschetz Trace Formula, 165, 214
- group sieve, 32, 48, 62, 70, 106, 115–117, 122, 125–127, 130, 141
- handlebody, 137, 141, 270
- Hasse inequality, 59
- Heegaard splitting, 137, 144, 281
- higher-dimensional large sieve, 48
- homotopy, 268, 271, 273
- hyperbolic lattice point problem, 102
- hyperbolic surfaces, 120
- hyperelliptic curve, 7, 160, 178, 185, 193, 201
- inclusion-exclusion principle, 45
- independence, 31, 254
- independent events, 47, 255
- independent random variables, 88, 108, 255
- induced representation, 73, 77, 84, 222, 226
- intersection pairing, 133, 141, 278
- invariant vector, 107, 108, 118, 226, 232, 233, 237, 239, 240, 242, 244
- irreducible character, 75, 224, 225
- irreducible polynomial, 52, 56, 205
- irreducible representation, 33, 34, 38, 70, 73, 84, 164, 182, 223–228, 234, 235
- isotopy, 277
- Jacobian variety, 7, 193, 201
- Jordan decomposition, 212
- Lagrangian subspace, xv, 141
- Lang–Kummer sheaf, 214
- large monodromy, 2, 183
- large sieve constant, 2, 9, 12, 18, 22, 25, 28, 33, 46, 52, 59, 63, 89–91, 97, 98, 106, 116, 123, 128, 130, 145, 164, 171
- large sieve inequality, 18, 91
- large sieve principle, 25
- larger sieve, 26
- Larsen alternative, 189
- lattice, 115, 234, 236
- lazy random walk, 90, 104
- left-invariant random walk, 6, 103, 106, 116, 122, 125, 127, 133, 145
- level of distribution, 51
- limit of discrete series, 235
- linear algebraic group, 72, 245–248
- linear group, 72, 74
- linear sieve, 200
- linearly disjoint, 31, 152, 164, 166, 173, 177, 179, 200
- lissee sheaf, 162, 163, 165, 168, 169
- loop, 112, 268–270, 275
- lower-bound sieve, 198, 200, 201, 203
- mapping class group, 6, 132–135, 137, 152, 268, 276, 277
- Markov process, 89
- matrix coefficient, 34, 35, 117, 223, 226
- maximal torus, 75, 212, 247
- mean-value theorem, 147, 196
- modular group, 101
- moduli space, 190
- monodromy group, 2, 154, 174
- Mordell–Weil group, 59
- multiplicative function, 27, 130, 181, 182, 262, 263
- multiplicative group, 64, 246, 247
- multiplicative large sieve inequality, 57
- multiplicator, xvi, 84, 85, 176, 177, 179, 180, 195, 210
- multiplicity, 41, 73, 75, 81, 117, 167, 225, 226
- naïve height, 60
- norm of a matrix, 102
- normal distribution, 92, 257
- numerator of the zeta function, 157, 159, 175–178, 183, 189, 192, 195
- orthogonality of characters, 38
- orthogonality relations, 118
- orthonormal basis, 12, 15, 23, 32, 34, 37, 224, 226
- periodicity, 90, 95, 104, 110, 120, 121, 137, 151
- Poincaré duality, 176
- pointwise pure, 166
- Poisson distribution, 88
- positive roots, 79
- positivity, 23, 50, 63, 64, 72, 73, 97, 110, 143
- presentation, 94, 114, 115, 151
- Prime Number Theorem, xv, 196

- prime sieve support, 9, 10, 18, 22, 29, 47, 63, 87, 97, 122, 125, 128, 172, 182, 195, 197, 200
- prime-to- p part, 74
- primitive character, 57, 58
- primitive irreducible representation, 42, 44
- primitive subspace, 11, 12, 14
- principal congruence subgroup, 102, 234
- principal divisor, 194
- principal series, 74, 77, 86, 228, 235, 236
- probabilistic sieve, 87, 94, 97, 104
- probability density, 10
- probability space, 87, 103, 254
- Property (τ), 104, 106, 109, 110, 112, 113, 115, 120, 140, 151, 232, 233, 236–238
- Property (T), 104, 115, 120, 124–126, 140, 151, 152, 232–238, 244
- pseudo-Anosov, 6, 132–136, 280–282
- radical, 78
- ramification, 170
- random 3-manifold, 94, 137–139, 145
- random group, 94, 95, 138
- random matrix, 96
- random products, 128
- random real number, 88
- random variable, 87, 254
- random walk, 5, 88, 95, 96, 121, 132, 148, 254, 259
- random walk on a graph, 111
- random walk on a group, 103, 105, 111, 116, 134
- random walks on $SL(n, \mathbf{Z})$, 145
- rank, 73, 82, 119, 212, 231, 247
- rational monodromy group, 178, 187, 189
- recurrent random walk, 259
- recurrent set, 94, 97, 144, 259, 261
- reducible characteristic polynomial, 132
- reducible polynomial, 94
- reductive group, 73, 75, 80–82, 187, 248
- regular character, 77, 85, 224
- regular element, 248
- regular representation, 76, 109, 224, 226
- relation of odd length, 113, 115, 120, 123, 125
- relative Property (T), 239
- representation theory, 32, 220, 223, 239
- residually finite, 151
- reversed characteristic polynomial, xvi, 163, 179, 182, 195, 211, 213, 215
- Riemann Hypothesis over finite fields, 130, 151, 154, 156, 166, 210
- rigidifying data, 190
- root subgroups, 114
- saving factor, 25, 26, 200, 210
- Selberg sieve, 197, 200
- Selberg's theorem, 81, 120, 236, 237
- self-adjoint operator, 109
- semisimple element, 211, 212, 248
- semisimple group, 115, 248
- semisimple rank, 78, 85, 249
- Siegel's theorem, 60
- sieve axioms, 13, 14
- sieve error term, 14
- sieve for Frobenius, 7, 36, 67, 154, 164, 168, 171, 175, 176, 179, 187, 191
- sieve of dimension zero, 68
- sieve problem, 9
- sieve setting, 8, 18, 21, 22, 25, 69, 87, 105, 197
- sieve support, 10, 18, 21, 26–28, 41, 46, 47, 52, 89, 91, 92, 97, 98, 116, 122, 125, 169, 172, 173, 175, 197
- sieving sets, 9, 23, 25, 26, 46, 47, 87, 130, 134, 191, 195, 199, 200, 202, 218
- siftable set, 8, 9, 18, 19, 21, 22, 30, 31, 34, 46, 48, 53, 60, 67, 69, 87, 89, 91, 97, 98, 102–105, 116, 122, 125, 126, 141, 161, 195, 197, 202
- sifted set, 9, 16, 87, 134, 197
- signed permutations, xvi, 127, 158
- simple random walk, 5
- simple random walk on \mathbf{Z} , 89, 92, 104
- simple random walk on \mathbf{Z}^s , 96, 259
- simply-connected algebraic group, 212, 237, 248
- simply-connected topological space, 269
- small sieve, 1, 197, 200, 263
- smoothing, 19
- special linear group, 113
- spectral gap, 81, 82, 112
- spectral radius, 108
- split algebraic group, 247
- split torus, 75, 77, 80, 81, 247
- splitting field, 52, 54, 157–160, 178–180, 185, 190, 192
- squarefree integer, 68, 88
- Steinberg character, 81
- strongly irreducible, 137
- sum of Betti numbers, 168
- sum of local traces of Frobenius, 165

- Swan conductor, 169–171
- symmetric generating set, 103, 105, 106, 113, 120, 122, 125, 127, 133, 135
- symplectic group, xv, 72, 74, 113, 246, 278, 279
- symplectic matrix, 102, 141, 182
- symplectic similitude, xvi, 84, 172, 176, 180, 195, 210, 246
- symplectic transvection, 279
- tamely ramified covering, 169, 170
- Tate twist, 166, 167
- Torelli group, 135, 279
- torus, 78, 247
- transient random walk, 94, 259
- transient set, 6, 98, 132, 133, 135, 136, 139, 145, 259
- trivial family, 177
- Turán's method, 24
- uniform density, 11
- uniform distribution, 112, 148
- uniformity, 3
- uniformity of sieves, 26, 145
- unimodular matrix, 102, 207
- unipotent element, 77, 212, 247
- unipotent group, 247–249
- unitary representation, 109, 115, 220, 232, 233, 235–238
- universal cover, 273
- upper-bound sieve, 198, 200
- variance, xv, 256
- variational characterization of eigenvalues, 109
- vector sieve, 1
- Virtual Haken Conjecture, 137, 138, 140
- Weierstrass equation, 5, 59
- weight, 166
- Weyl criterion, 201
- Weyl group, 73, 79, 127, 131
- word-length metric, 103
- words, 103
- Zariski closure, 178, 187
- zeta function, 2, 156, 185, 194, 266