Interactions of Computational Complexity Theory and Mathematics

Avi Wigderson

October 22, 2017

Abstract

[This paper is a (self contained) chapter in a new book on computational complexity theory, called *Mathematics and Computation*, whose draft is available at https://www.math.ias.edu/avi/book].

We survey some concrete interaction areas between computational complexity theory and different fields of mathematics. We hope to demonstrate here that hardly any area of modern mathematics is untouched by the computational connection (which in some cases is completely natural and in others may seem quite surprising). In my view, the breadth, depth, beauty and novelty of these connections is inspiring, and speaks to a great potential of future interactions (which indeed, are quickly expanding). We aim for *variety*. We give short, simple descriptions (without proofs or much technical detail) of ideas, motivations, results and connections; this will hopefully entice the reader to dig deeper. Each vignette focuses only on a single topic within a large mathematical filed. We cover the following:

- Number Theory: *Primality testing*
- Combinatorial Geometry: *Point-line incidences*
- Operator Theory: *The Kadison-Singer problem*
- Metric Geometry: *Distortion of embeddings*
- Group Theory: *Generation and random generation*
- Statistical Physics: *Monte-Carlo Markov chains*
- Analysis and Probability: *Noise stability*
- Lattice Theory: *Short vectors*
- Invariant Theory: *Actions on matrix tuples*

1 introduction

The Theory of Computation (ToC) lays out the mathematical foundations of computer science. I am often asked if ToC is a branch of Mathematics, or of Computer Science. The answer is easy: it is clearly both (and in fact, much more). Ever since Turing's 1936 definition of the *Turing machine*, we have had a formal mathematical model of computation that enables the rigorous mathematical study of computational tasks, algorithms to solve them, and the resources these require. At the same time, the simple description of the Turing machine allowed its simple logical structure to be implemented in hardware, and its universal applicability fueled the rapid development of computer technology, which now dominates our life.

Computation was part mathematics from its origins, and motivated many of its developments. Algorithmic questions have occupied mathematicians throughout history (as elaborated in the introduction to [Wig17]), and this naturally grew considerably when computers arrived. However, the advent of *computational complexity theory* over the past few decades has greatly expanded and deepened these connections. The study of new diverse models generated and studied in complexity theory broadened the nature of mathematical problems it encountered and formulated, and the mathematical areas and tools which bear upon these problems. This expansion has led to numerous new interactions that enrich both disciplines. This survey tells the stories of some of these interactions with different mathematical fields, illustrating their diversity.

We note in passing that a similar explosion of connections and interactions is underway between ToC and practically *all* sciences. These stem from computational aspects of diverse natural processes, which beg for algorithmic modeling and analysis. As with mathematics, these interactions of ToC with the sciences enrich both sides, expose *computation* as a central notion of intellectual thought, and highlights its study as an independent discipline, whose mission and goals expand way beyond those emanating from its parent fields of Math and CS. But this is the subject of a different survey (which I partly provide in the last chapter of [Wig17]).

Back to the interactions of computational complexity theory and different areas of math. I have chosen to focus on essentially one problem or development within each mathematical field. Typically this touches only a small subarea, and does not do justice to a wealth of other connections. Thus each vignette should be viewed as a demonstration of a larger body of work and even bigger potential. Indeed, while in some areas the collaborations are quite well established, in others they are just budding, with lots of exciting problems waiting to be solved and theories to be developed. Furthermore, the connections to algorithms and complexity (which I explain in each) are quite natural in some areas, but quite surprising in others. While the descriptions of each topic are relatively short, they include background and intuition, as well as further reading material. Indeed, I hope these vignettes will tempt the reader to explore further.

Here is a list of the covered areas and topics chosen in each; these sections can be read in any order. The selection of fields and foci is affected by my personal taste and limited knowledge. More connections to other fields like Combinatorics, Optimization, Logic, Topology and Information Theory appear in parts of the book [Wig17].

- Number Theory: *Primality testing*
- Combinatorial Geometry: *Point-line incidences*
- Operator Theory: *The Kadison-Singer problem*
- Metric Geometry: *Distortion of embeddings*
- Group Theory: *Generation and random generation*
- Statistical Physics: *Monte-Carlo Markov chains*
- Analysis and Probability: *Noise stability*
- Lattice Theory: *Short vectors*
- Invariant Theory: *Actions on matrix tuples*

2 Number Theory

As mentioned, the need to efficiently compute mathematical objects has been central to mathematicians and scientists throughout history, and of course the earliest subject is arithmetic. Perhaps the most radical demonstration is the place value system we use to represent integers, which is in place for Millenia precisely due to the fact that it supports extremely efficient manipulation of arithmetic operations. The next computational challenge in arithmetic, since antiquity, was accessing the multiplicative structure of integers represented this way.

Here is an except from C. F. Gauss' appeal[1] to the mathematics community of his time (in article 329 of *Disquisitiones Arithmeticae* (1801)), regarding the computational complexity of *testing primality* and *integer factorization*. The importance Gauss assigns to this computational challenge, his frustration of the state of art, and his imploring the mathematical community to resolve it shine through!

The problem of distinguishing prime numbers from composite numbers, and of resolving the latter into their prime factors is known to be one of the most important and useful in arithmetic. It has engaged the industry and wisdom of ancient and modern geometers to such an extent that it would be superfluous to discuss the problem at length. Nevertheless we must confess that all methods that have been proposed thus far are either restricted to very special cases or are so laborious and difficult that even for numbers that do not exceed the limits of tables constructed by estimable men, they try the patience of even the practiced calculator. And these methods do not apply at all to larger numbers ... the dignity of the science itself seems to require that every possible means be explored for the solution of a problem so elegant and so celebrated.

We briefly recount the state-of-art of these two basic algorithmic problems in number theory. A remarkable response to Gauss' first question, *efficiently deciding primality*, was found in 2002 by Agrawal, Kayal, and Saxena [AKS04]. The use of symbolic polynomials for this problem is completely novel. Here is their elegant characterization of prime numbers.

Theorem 2.1 [AKS04] *An integer $N \geq 2$ is prime if and only if*

- *N is not a perfect power,*
- *N does not have any prime factor $\leq (\log N)^4$,*
- *For every $r, a < (\log N)^4$ we have the following equivalence of polynomials over $\mathbb{Z}_N[X]$:*

$$(X + a)^N \equiv X^N + a \mod (X^r - 1)$$

It is not hard to see that this characterization gives rise to a simple algorithm for testing primality that is deterministic, and runs in time that is *polynomial* in the binary description length of N. Previous deterministic algorithms either assumed the generalize Riemann hypothesis [Mil76] or required slightly superpolynomial time [APR83]. The AKS deterministic algorithm came after a sequence of efficient *probabilistic* algorithms [SS77, Rab80, GK86, AH92], some elementary and some requiring sophisticated use and development of number theoretic techniques. These probabilistic and deterministic algorithms were partly motivated by, and are important to the field of cryptography.

What is not so well-known, even for those who did read the beautiful, ingenious proof in [AKS04], is that AKS developed their deterministic algorithm by carefully "de-randomizing" a previous probabilistic algorithm for primality of [AB03] (which uses polynomials). We note that *de-randomization*, the conversion of probabilistic algorithms into deterministic ones, is by now a major area in computational complexity with a rich theory, and many other similar successes as well as challenges. The stunning possibility that *every* efficient probabilistic algorithm has a deterministic counterpart is one of the major problems of computational complexity, and there is strong evidence supporting it (see [IW97]). Much more on this can be found in the randomness chapters of [Wig17].

Gauss' second challenge, of whether efficiently factoring integers is possible, remains open. But this very challenge has enriched computer science, both practical and theoretical in several major ways. Indeed, the

[1] Which is of course in Latin. I copied this English translation from a wonderful survey of Granville [Gra05] on the subject matter of this section.

assumed hardness of factoring is the main guarantee of security in almost all cryptographic and e-commerce systems around the world (showing that difficult problems can be useful!). More generally, cryptography is an avid consumer of number theoretic notions, including elliptic curves, Weil pairings, and more, which are critical to a variety of cryptographic primitives and applications. These developments shatter Hardy's view of number theory as a completely useless intellectual endeavor.

There are several problems on integers whose natural definitions depend on factorization, but can nevertheless be solved efficiently, bypassing the seeming need to factor. Perhaps the earliest algorithm ever formally described is Euclid's algorithm for computing the GCD (greatest common divisor) of two given integers[2] m and n. Another famous such algorithm is for computing the Legendre-Jacobi symbol $(\frac{m}{n})$ via Gauss' law of quadratic reciprocity.

A fast algorithm for factoring may come out of left-field with the new development of quantum computing, the study of computers based on quantum-mechanical principles, which we discussed in the quantum chapter of the book [Wig17]. Shor has shown in [Sho94] that such computers are capable of factoring integers in polynomial time. This result led governments, companies, and academia to invest billions in developing technologies which will enable building large-scale quantum computers, and the jury is still out on the feasibility of this project. There is no known theoretical impediment for doing so, but one possible reason for failure of this project is the existence of yet-undiscovered principles of quantum mechanics.

Other central computational problems include solving polynomial equations in finite fields, for which one of the earliest efficient (probabilistic) algorithm was developed by Berlekamp [Ber67] (it remains a great challenge to de-randomize this algorithm!). Many other examples can be found in the Algorithmic Number Theory book [BS97].

3 Combinatorial geometry

What is the smallest area of a planar region which contains a unit length segment in *every* direction? This is the Kakeya needle problem (and such sets are called *Kakeya sets*), which was solved surprisingly by Besicovich [Bes19] who showed that this area can be arbitrarily close to zero! Slight variation on his method produces a Kakeya set of Lebesque measure zero. It makes sense to replace "area" (namely, Lesbegue measure) by the more robust measures, such as the Hausdorff and Minkowski dimensions. This changes the picture: Davies [Dav71] proved that a Kakeya set in the plane must have full dimension (=2) in both measures, despite being so sparse in Lebesgue measure.

It is natural to extend this problem to higher dimensions. However, obtaining analogous results (namely, that the Hausdorff and Minkowski dimensions are full) turns out to be extremely difficult. Despite the seemingly recreational flavor, this problem has significant importance in a number of mathematical areas (Fourier analysis, Wave equations, analytic number theory, and randomness extraction), and has been attacked through a considerable diversity of mathematical ideas (see [Tao09]).

The following finite field analogue of the above Euclidean problem was suggested by Wolff [Wol99]. Let $\mathbb{F}$ denote a finite field of size q. A set $K \subseteq \mathbb{F}^n$ is called Kakeya if it contains a line in every direction. More precisely, for every direction $b \in \mathbb{F}^n$ there is a point $a \in \mathbb{F}^n$ such that the line $\{a + bt : t \in \mathbb{F}\}$ is contained in K. As above, we would like to show that any such K must be large (think of the dimension n as a large constant, and the field size q as going to infinity).

Conjecture 3.1. Let $K \subseteq \mathbb{F}^n$ be a Kakeya set. Then $|K| \geq C_n q^n$, where C_n is a constant depending only on the dimension n.

The best exponent of q in such a lower bound intuitively corresponds to the Hausdorff and Minkowski dimensions in the Euclidean setting. Using sophisticated techniques from arithmetic combinatorics, Bourgain, Tao and others improved the trivial bound of $n/2$ to about $4n/7$.

Curiously, the exact same conjecture arose, completely independently, within ToC, from the work [LRVW03] on *randomness extractors*, an area which studies the "purification" of "weak random sources" (see e.g. the

[2] It extends to polynomials, and allows efficient way of computing multiplicative inverses in quotient rings of $\mathbb{Z}$ and $\mathbb{F}[x]$.

survey [Vad11] on this important notion). In [LRVW03] Wolff's conjecture takes a probabilistic form, asking about the (min)-entropy of a random point on a random line in a Kakeya set. With this motivation, Dvir [Dvi09] brilliantly proved the Wolff conjecture (sometimes called the Finite Field Kakeya conjecture), using the (algebraic-geometric) "polynomial method" (which is inspired by techniques in decoding algebraic error-correcting codes). Many other applications of this technique to other geometric problems quickly followed, including the Guth-Katz [GK10] resolution of the famous Erdős distance problem, as well as for optimal randomness extraction and more (some are listed in Dvir's survey [Dvi10]).

Subsequent work determined the exact value of the constant C_n above (up to a factor of 2) [DKSS13].

Theorem 3.2 [DKSS13] *Let $K \subseteq \mathbb{F}^n$ be a Kakeya set. Then $|K| \geq (q/2)^n$. On the other hand, there exist Kakeya sets of size $\leq 2 \cdot (q/2)^n$.*

Many other problems regarding incidences of points and lines (and higher-dimensional geometric objects) have been the source of much activity and collaboration between geometers, algebraists, combinatorialists and computer scientists. The motivation for these questions in the computer science side come from various sources, e.g. problems on local correction of errors [BDWY13] and de-randomization [DS07, KS09]. Other incidence theorems, e.g. Szemerédi-Trotter [STJ83] and its finite field version of Bourgain-Katz-Tao [BKT04] have been used e.g. in randomness extraction [BIW06] and compressed sensing [GLR10].

4 Operator theory

The following basic mathematical problem of Kadison and Singer from 1959 [KS59] was intended to formalize a basic question of Dirac concerning the "universality" of measurements in quantum mechanics. We need a few definitions. Consider $B(\mathcal{H})$, the algebra of continuous linear operators on a Hilbert space $\mathcal{H}$. Define a *state* to be a linear functional f on $B(\mathcal{H})$, normalized to $f(I) = 1$, which takes non-negative values on positive semidefinite operators. The states form a convex set, and a state is called *pure* if it is not a convex combination of other states. Finally, let D be the sub-algebra of $B(\mathcal{H})$ consisting of all *diagonal* operators (after fixing some basis).

Kadison and Singer asked if every pure state on D has a *unique* extension to $B(\mathcal{H})$. This problem on infinite-dimensional operators found a host of equivalent formulations in finite dimensions, with motivations and intuitions from operator theory, discrepancy theory, Banach space theory, signal processing, and probability. All of them were solved affirmatively in recent work of Marcus, Spielman, and Srivastava [MSS13b] (which also surveys the many related conjectures). Here is one statement they prove, which implies the others.

Theorem 4.1 [MSS13b] *For every $\epsilon > 0$, there is an integer $k = k(\epsilon)$ so that the following holds. Fix any n and any $n \times n$ matrix A with zeros on the diagonal and of spectral norm 1. Then there is a partition of $\{1, 2, \cdots, n\}$ into k subsets, $S_1, S_2, \cdots, S_k$, so that each of the principal minors A_i (namely A restricted to rows and columns in S_i) has spectral norm at most ϵ.*

This statement clearly implies that one of the minors has linear size, at least n/k. This consequence is known as the *Restricted Invertibility* Theorem of Bourgain and Tzafriri [BT91], itself an important result in operator theory.

How did computer scientists get interested in this problem? Without getting into too many details, here is a sketchy description of the meandering path which led to this spectacular result.

A central computational problem, at the heart of numerous applications, is solving a linear system of equations. While Gaussian elimination does the job quite efficiently (the number of arithmetic operations is about n^3 for $n \times n$ matrices), for large n this is still inefficient. Thus faster methods are sought, hopefully nearly linear in the number of non-zero entries of the given matrix. For *Laplacian*[3] linear systems (arising in many graph theory applications, such as computing electrical flows and random walks), Spielman and Teng [ST11] achieved precisely that! A major notion they introduced was *spectral sparsifiers* of matrices (or equivalently, weighted graphs).

[3] Simply, symmetric PSD matrices with zero row sum.

A sparsifier of a given matrix is another matrix, with far fewer (indeed, linear) non-zero entries, which nevertheless has essentially the same (normalized) spectrum as the original (it is not even obvious that such a sparse matrix exists). We note that a very special case of sparsifiers of complete graphs are by definition *expander graphs*[4] (see much more about this central concept of expanders in [HLW06, Wig17]). The algorithmic applications led to a quest for optimal constructions of sparsifiers for arbitrary Laplacian matrices (in terms of trade-off between sparsity and approximation), and these were beautifully achieved in [BSS14] (who also provided a deterministic polynomial time algorithm to construct such sparsifiers). This in turn has led [SS12] to a new proof, with better analysis, of the Restricted Invertibility theorem mentioned above, making the connection to the Kadison-Singer problem.

However, the solution to Kadison-Singer seemed to require another detour. The same team [MSS13a] first resolved a bold conjecture of Bilu and Linial [BL06] on the spectrum of "signings" of matrices[5]. This conjecture was part of a plan for a *simple*, iterative construction of Ramanujan graphs, the best[6] possible expander graphs. Ramanujan graphs were introduced and constructed in [LPS88, Mar88], but rely on deep results in number theory and algebraic geometry (believed by some to be essential for *any* such construction). Bilu and Linial sought instead an elementary construction, and made progress on their conjecture, showing how their iterative approach gives yet another way to construct "close to" Ramanujan expanders.

To prove the Bilu-Linial conjecture (and indeed produce bipartite Ramanujan graphs of every possible degree—something the algebraic constructions couldn't provide), [MSS13a] developed a theory of *interlacing polynomials* that turned out to be the key technical tool for resolving Kadison-Singer in [MSS13b]. In both cases, the novel view is to think of these conjectures probabilistically, and analyze the norm of a random operator by analyzing the average characteristic polynomial. That this method makes sense and actually works is deep and mysterious. Moreover, it provides a new kind of existence proofs for which no efficient algorithm (even probabilistic) of finding the desired objects is known. The analysis makes heavy use of the theory of *Real stable* polynomials, and the inductive process underlying it is reminiscent (and inspired by) Gurvits' [Gur08] remarkable proof of the van der Waerden conjecture and its generalizations[7].

5 Metric Geometry

How close one metric space is to another is captured by the notion of *distortion*, measuring how distorted distances of one become when embedded into the other. More precisely,

Definition 5.1. Let (X, d) and (X', d') be two metric spaces. An embedding $f : X \to X'$ has distortion $\leq c$ if for every pair of points $x, y \in X$ we have

$$d(x, y) \leq d'(f(x), f(y)) \leq c \cdot d(x, y).$$

When X is finite and of size n, we allow $c = c(n)$ to depend on n.

Understanding the best embeddings between various metric and normed spaces has been a long endeavor in Banach space theory and metric geometry. An example of one major result in this area is Bourgain's embedding theorem [Bou85].

Theorem 5.2 [Bou85] *Every metric space of size n can be embedded into Euclidean space L_2 with distortion $O(\log n)$.*

The first connection between these structural questions and computational complexity was made in the important paper of Linial, London and Rabinovich [LLR95]. They asked for efficient algorithms for actually

[4] All non-trivial eigenvalues of the complete graph (or constant matrix) are 0, and an expander is a sparse graph in which all non-trivial eigenvalues are tiny.

[5] Simply, this beautiful conjecture states that for *every* d-regular graph, there exist $\{-1, 1\}$ signs of the edges which make all eigenvalues of the resulting signed adjacency matrix lie in the "Ramanujan interval" $[-2\sqrt{d-1}, 2\sqrt{d-1}]$.

[6] With respect to the spectral gap. This is one of a few important expansion parameters to optimize.

[7] This is yet another example of structural result (on doubly stochastic matrices) whose proof was partly motivated by algorithmic ideas. The connection is the use of hyperbolic polynomials in optimization (more specifically, as barrier functions in interior point methods.

finding embeddings of low distortion, and noticed that for some such problems it is natural to use semi-definite programming. They applied this geometric connection to get old and new results for algorithmic problems on graphs (in particular, the sparsest cut problem we will soon discuss. Another motivation they discuss (which quickly developed into a major direction in approximation algorithms) is that some computations (e.g. finding nearest neighbors) are more efficient in some spaces than others, and so *efficient*, low-distortion embedding may provide useful reductions from harder to easier space. They describe such an efficient algorithm implementing Bourgain's Theorem 5.2 above, and also prove that his bound is best possible (the metric proving it is simply the distances between points in any constant-degree *expander* graph[8]).

The next shift in the evolution of this field, and in the level of interactions between geometers and ToC researchers, came from trying to prove "hardness of approximation" results. One example is the Goemans-Linial conjecture [Goe97, Lin02], studying the sparsest cut problem, about the relation between L_1 and the "negative type" metric space L_2^2 (a general class of metrics which arise naturally in several contexts). Roughly, thes are metrics on $\mathbb{R}^n$ in which Euclidean distances are squared. More precisely, a metric (X, d) is of negative type (namely, in L_2^2), if $(X, \sqrt{d})$, is isometric (has no distortion) to a subset of L_2.

Conjecture 5.3. Every L_2^2 metric can be embedded into L_1 with constant distortion.

This conjecture was proved false by Khot and Vishnoi [KV05]:

Theorem 5.4 [KV05] *For every n there are n-point subsets of L_2^2 for which every embedding to L_1 requires distortion $\Omega(\log\log n)^{1/6}$.*

Far more interesting than the result itself is its origin. Khot and Vishnoi were trying to prove that the (weighted) "sparsest cut" problem is hard to approximate. They managed to do so under a computational assumption, known as the *Unique Games* conjecture of Khot [Kho02] via a so-called *PCP*-reduction (see also [Kho10, Wig17]). The elimination of this computational assumption is the magical part, that demonstrates the power and versatility of reductions between computational problems. They apply their PCP reduction to a *particular*, carefully chosen unique games instance, which cannot be well approximated by a certain semi-definite program. The outcome was an instance of the sparsest cut problem which the same reduction ensures is hard to approximate by a semi-definite program. As discussed above, that outcome instance could be understood as a metric space, and the hardness of approximation translates to the required distortion bound!

The exact distortion of embedding L_2^2 into L_1 has been determined precisely to be $\sqrt{\log n}$ (up to lower order factors) in two beautiful sequences of works developing new algorithmic and geometric tools; we mention only the final word for each, as these papers contain a detailed history. On the upper bound side, the efficient algorithm approximating non-uniform sparsest cut to a factor $\sqrt{\log n}\log\log n$, which yields the same distortion bound, was obtained by Arora, Lee and Naor [ALN08] via a combination of the so-called "chaining argument" of [SSU04] and the "measured descent" embedding method of [KLMN05]. A lower bound of $\sqrt{\log n}$ on the distortion was very recently proved by Naor and Young [NY17] using a new isoperimetric inequality on the Heisenberg group.

Another powerful connection between such questions and ToC is through (again) expander graphs. A basic example is that the graph metric of any constant-degree expander proves that Bourgain's embedding theorem above is optimal! Much more sophisticated examples arise from trying to understand (and perhaps disprove) the Novikov and the Baum-Connes conjectures (see [KY06]). This program relies on another, much weaker notion of *coarse* embedding.

Definition 5.5. (X, d) has a coarse embedding into (X', d') if there is a map $f : X \to X'$ and two increasing, unbounded real functions α, β such that for every two points $x, y \in X$,

$$\alpha(d(x, y)) \leq d'(f(x), f(y)) \leq \beta(d(x, y)).$$

Gromov [Gro87] was the first to construct a metric (the word metric of a group) which cannot be coarsely embedded into a Hilbert space. His construction uses an infinite family of *Cayley* expanders (graphs

[8]The presence of such graphs in different sections illustrate how fundamental they are in diverse mathematical areas, and the same holds for algorithms and complexity theory.

defined by groups). This result was greatly generalized by Lafforgue [Laf08] and Mendel-Naor [MN14], who constructed graph metrics that cannot be coarsely embedded into any *uniformly convex* space. It is interesting that while Lafforgue's method is algebraic, the Mendel-Naor construction follows the combinatorial *zig-zag* construction of expanders [RVW02] from computational complexity.

Many other interaction projects regarding metric embeddings and distortion we did not touch on include their use in numerous algorithmic and data structure problems like clustering, distance oracles the k-server problem, as well as the fundamental interplay between distortion and *dimension reduction* relevant to both geometry and CS, where so many basic problems are open.

6 Group Theory

Group theorists, much like number theorists, have been intrinsically interested in computational problems since the origin of the field. For example, the *word problem* (given a word in the generators of some group, does it evaluate to the trivial element?) is so fundamental to understanding any group one studies, that as soon as language was created to formally discuss the computational complexity of this problem, hosts of results followed trying to pinpoint that complexity. These include decidability and undecidability results once Turing set up the theory of computation and provided the first undecidable problems, and these were followed with $\mathcal{NP}$-completeness results and efficient algorithms once $\mathcal{P}$ and $\mathcal{NP}$ were introduced around 1970. Needless to say, these *algorithmic results* inform of *structural* complexity of the groups at hand. And the word problem is but the first example. Another demonstration is the beautiful interplay between algorithmic and structural advances over decades, on the *graph isomorphism problem*, recently leading to breakthrough of Babai [Bab15]! A huge body of work is devoted to finding efficient algorithms for computing commutator subgroups, Sylow subgroups, centralizers, bases, representations, characters, and a host of other important substructures of a group from some natural description of it. Excellent textbooks include [HEO05, Ser03].

Here we focus on two related problems, the *generation* and *random generation* problems, and new conceptual notions borrowed from computational complexity which are essential for studying them. Before defining them formally (below), let us consider an example. Assume I hand you 10 invertible matrices, say 100×100 in size, over the field of size 3. Can you tell me if they generate another such given matrix? Can you even produce convincing evidence of this before we both perish? How about generating a random matrix in the subgroup spanned by these generators? The problem, of course, is that this subgroup will have size far larger than the number of atoms in the known universe, so its elements cannot be listed, and typical words generating elements in the group may need to be prohibitively long. Indeed, even the extremely special cases, for elements in $\mathbb{Z}_p^*$ (namely one, 1×1 matrix), the first question is related to the *discrete logarithm* problem, and for $\mathbb{Z}_{p \cdot q}^*$ it is related to the *integer factoring* problem, both currently requiring exponential time to solve (as a function of the description length).

Let us consider any finite group G and let $n \approx \log |G|$ be roughly the length of a description of an element of G. Assume we are given k elements in G, $S = \{s_1, s_2, \ldots, s_k\}$. It would be ideal if the procedures we describe would work in time polynomial in n and k (which prohibits enumerating the elements of G, whose size is exponential in n).

The *generation problem* asks if a given element $g \in G$ is generated by S. How does one prove such a fact? A standard certificate for a positive answer is a *word* in the elements of S (and their inverses) which evaluates to g. However, even if G is cyclic, the shortest such word may be exponential in n. An alternative, computationally motivated description, is to give a *program* for g. Its definition shows that the term "program" suits it perfectly, as it has the same structure as usual computer programs, only that instead of applying some standard Boolean or arithmetic operations, we use the group operations of multiplication and inverse.

Definition 6.1. A *program* (over S) is a finite sequence of elements $g_1, g_2, \cdots, g_m$, where every element g_i is either in S, or is the inverse of a previous g_j, or is the product of previous g_j, g_ℓ. We say that it computes g simply if $g = g_m$.

In the cyclic case, programs afford exponential savings over words in description length, as a program

allows us to write large powers by repeatedly squaring elements. What is remarkable is that such savings are possible for *every* group. This discovery of Babai and Szemerédi [BS84] says that every element of every group has an extremely succinct description in terms of any set of elements generating it.

Theorem 6.2 [BS84] *For every group G, if a subset of elements S generates another element g, then there is a program of length at most $n^2 \approx (\log |G|)^2$ which computes g from S.*

It is interesting to note that the proof uses a structure which is very combinatorial and counterintuitive for group theorists: that of a *cube*, which we will see again later. For a sequence $(h_1, h_2, \cdots, h_t)$ of elements from G, the cube $C(h_1, h_2, \cdots, h_t)$ is the (multi)set of 2^t elements $\{h_1^{\epsilon_1}, h_2^{\epsilon_2}, \cdots, h_t^{\epsilon_t}\}$, with $\epsilon_i \in \{0, 1\}$. Another important feature of the proof is that it works in a very general setting of "black-box" groups—it never needs an explicit description of the host group, only the ability to multiply elements and take their inverses. This is a very important paradigm for arguing about groups, and will be used again below.

How does one prove that an element g is *not* generated by S? It is possible that there is no short "classical" proof! This question motivated Babai to define Arthur-Merlin games—a new notion of probabilistic, interactive proofs (simultaneously with Goldwasser, Micali, and Rackoff [GMR89], who proposed a similar notion for cryptographic reasons), and showed how non-membership can be certified in this new framework. The impact of the definition of interactive proofs on the theory of computation has been immense, and is discussed in e.g. in the books [Gol08, AB09, Wig17].

Returning to the generation problem, let us now consider the problem of *random generation*. Here we are given S, and would like a randomized procedure which will quickly output an (almost) uniform distribution on the subgroup H of G generated by S. This problem, besides its natural appeal, is often faced by computational group theorists, being a subroutine in many group-theoretic algorithms. In practice often heuristics are used, like the famous "product replacement algorithm" and its variants, which often work well in practice (see e.g. the recent [BLG12] and references). We will discuss here provable bounds.

It is clear that sufficiently long random words in the elements of S and its inverses will do the job, but just as with certificates, sufficiently long is often prohibitively long. In a beautiful paper, Babai [Bab91] describes a certain process generating a random program which computes a nearly-uniform element of H, and runs in time $n^5 \approx (\log |G|)^5$ steps. It again uses cubes, and works in the full generality of black-box groups. This paper was followed by even faster algorithms with simpler analysis by Cooperman and by Dixon [Coo02, Dix08], and the state-of-art is an algorithm whose number of steps is remarkably the same as the length of proofs of generation above—in other words, randomness roughly achieves the efficiency of non-determinism for this problem. Summarizing:

Theorem 6.3 [Bab91, Coo02, Dix08] *For every group G, there is a* probabilistic *program of length* $\mathrm{poly}(n) \approx \mathrm{poly}(\log |G|)$ *that, given* any *generating set S for G, produces with high probability a (nearly) uniformly random element of G.*

7 Statistical Physics

The field of statistical physics is huge, and we focus here mainly on connections of statistical mechanics with the theory of computation. Numerous mathematical models exist of various physical and chemical systems, designed to understand basic properties of different materials and the dynamics of basic processes. These include such familiar models as Ising, Potts, Monomer-Dimer, Spin-Glass, Percolation, etc. A typical example explaining the connection of such mathematical models to physics and chemistry, and the basic problems studied is the seminal paper of Heilmann and Lieb [HL72].

Many of the problems studied can be viewed in the following general setting. We have a huge (exponential) space of objects called Ω (these objects may be viewed as the different configurations of a system). Each object is assigned a nonnegative weight (which may be viewed as the "energy" of that state). Scaling these weights gives rise to a probability distribution (often called the Gibbs distribution) on Ω, and to study its properties (phase transitions, critical temperatures, free energy, etc.) one attempts to generate samples from this distribution. Note that if the description of a state takes n bits, then brute-force listing of all probabilities in question is exponentially prohibitive. Thus efficiency of the sampling procedure is essential

to this study.

As Ω may be highly unstructured, the most common approach to this sampling problem is known as "Monte Carlo Markov Chain" (or "MCMC") method. The idea is to build a graph on the objects of Ω, with a pair of objects connected by an edge if they are similar in some sense (e.g. sequences which differ only in a few coordinates). Next, one starts from any object, and performs a biased random walk on this graph for some time, and the object reached is the sample produced. In many settings it is not hard to set up the random walk (often called Glauber dynamics or the Metropolis algorithm) so that the *limiting* distribution of the Markov chain is indeed the desired distribution. The main question in this approach is *when* to stop the walk and output a sample; *when* are we close enough to the limit? In other words, how long does it take the chain to converge to the limit? In most cases, these decisions were taken on intuitive, heuristic grounds, without rigorous analysis of convergence time. The exceptions where rigorous bounds were known were typically structured, e.g. where the chain was a Cayley graph of a group (e.g. [Ald83, Dia88]).

This state of affairs has changed considerably since the interaction in the past couple of decades with the theory of computation. Before describing it, let us see where computational problems even arise in this field. The two major sources are *optimization* and *counting*. That the setting above suits many instances of optimization problems is easy to see. Think of Ω as the set of solutions to a given optimization problem (e.g. the values of certain parameters designed to satisfy a set of constraints), and the weights representing the quality of a solution (e.g. the number of constraints satisfied). So, picking at random from the associated distribution favors high-quality solutions. The counting connection is more subtle. Here Ω represents a set of combinatorial objects one wants to count or approximate (e.g. the set of perfect matchings in a graph, or satisfying assignments to a set of constraints). It turns out that for very general situations of this type, sampling an object (approximately) at random is tightly connected to counting their number; it often allows a recursive procedure to approximate the size of the set [JVV86]. An additional observation is that viewing a finite set as a fine discretization of a continuous object (e.g. fine lattice points in a convex set) allows one to compute volumes and more generally integrate functions over such domains.

Around 1990, rigorous techniques were introduced [Ald90, Bro89, SJ89, DFK91] to analyze the convergence rates of such general Markov chains arising from different approximation algorithms. They establish *conductance* bounds on the Markov chains, mainly via *canonical paths* or *coupling* arguments (a survey of this early work is [JS96]). Collaborative work was soon able to formally justify the physical intuition behind some of the suggested heuristics for many models, and moreover drew physicists to suggest such ingenious chains for optimization problems. The field drew in probabilists and geometers as well, and by now is highly active and diverse. We mention two results to illustrate rigorous convergence bounds for important problems of this type.

Theorem 7.1 [JSV04] *The permanent of any nonnegative $n \times n$ matrix can be approximated, to any multiplicative factor $(1+\epsilon)$, in polynomial time in n/ϵ.*

The importance of this approximation algorithm stems from the seminal result of Valiant [Val79] about the permanent polynomial (that notorious sibling of the determinant polynomial, that looks identical except that the permanent has no signs; for more see [SY10, Wig17]). Valiant proved that the permanent is *universal*, capturing (via efficient reductions) essentially all natural counting problems, including those arising in the statistical physics models and optimization and counting problems above. So, unlike determinant, computing the permanent *exactly* is extremely difficult (harder than $\mathcal{NP}$-complete).

Theorem 7.2 [DFK91] *The volume of any convex set in n dimensions can be approximated, to any multiplicative factor $(1+\epsilon)$, in polynomial time in n/ϵ.*

The volume, besides its intrinsic interest, captures as well natural counting problems, e.g. the number of linear extensions of a given partially ordered set. The analysis of this algorithm, as well as its many subsequent improvements has used and developed purely structural results of independent interest in differential and convex geometry. It also led to generalizations, like efficiently sampling from any log-concave distribution (see the survey [Vem05]).

Another consequence of this collaboration was a deeper understanding of the relation between *spacial* properties (such as phase transitions, and long-range correlations between distant sites in the Gibbs distri-

bution) and *temporal* properties (such as speed of convergence of the sampling or approximately counting algorithms, like Glauber dynamics). This connection (surveyed e.g. in [DSVW04]) was established by physicists for spin systems since the 1970s. The breakthrough work of Weitz [Wei06] on the *hard core* model gave an *deterministic* algorithm which is efficient up to the phase transition, and this was complemented by a hardness result of Sly [Sly10] beyond the phase transition. These phase transition of computational complexity, at the same point as the phase transition of the Gibbs distribution are striking, and the generality of this phenomenon is still investigated.

More generally, the close similarity between statistical physics models and optimization problems, especially on random instances, is benefitting both sides. Let us mention a few exciting developments. It has unraveled the fine geometric structure of the space of solutions at the phase transition, pinpointing it e.g. for k-SAT in [ACORT11]. At the same time, physics intuition based on such ideas as renormalization, annealing, and replica symmetry breaking, has led to new algorithms for optimization problems, some of them now rigorously analyzed, e.g. as in [JS93]. Others, like one of the fastest (yet unproven) heuristics for such problems as Boolean Satisfiability (which is $\mathcal{NP}$-complete in general) are based on the physics method of "survey propagation" of [MPZ02]. Finally, new algorithmic techniques for similar physics and optimization problems, originate from an unexpected source, the *Lovasz Local Lemma* (LLL). The LLL is a probabilistic proof technique for the existence rare events in a probability space. Its efficient versions, formulating it algorithmically as a *directed, non-reversible* Markov chains, starting with the works of Moser [Mos09, MT10], have led to approximate counting and sampling versions for such events (see e.g. [GJL16]). A completely different, *deterministic* algorithm of Moitra [Moi16] for the LLL regime (of rare events) promises many more applications: it works even when the solution space (and hence the natural Markov chain) is not connected!

8 Analysis and Probability

This section gives a taste of a growing number of families of inequalities—large deviation inequalities, isoperimetric inequalities, etc.—that have been generalized beyond their classical origins due to a variety of motivations in the theory of computing and discrete mathematics. Further, the applications sometimes call for *stability* versions of these inequalities, namely an understanding of the structures which make an inequality nearly sharp. Here too these motivations pushed for generalizations of classical results and many new ones. Most of the material below, and much more on the motivations, applications and developments in this exciting area of the analysis of Boolean functions, can be found in the book [O'D14] by O'Donnell.

The following story can be told from several angles. One is the *noise sensitivity* of functions. We restrict ourselves to the Boolean cube endowed with the uniform probability measure, but many of the questions and results extend to arbitrary product probability spaces. Let $f : \{-1, 1\}^n \to \mathbb{R}$, which we assume is balanced, namely $E[f] = 0$. When the image of f is $\{-1, 1\}$, we can think of f as a voting scheme, translating the binary votes of n individuals into a binary outcome. One natural desire from such a voting scheme may be *noise stability*—that typically very similar inputs (vote vectors) will yield the same outcome. While natural in this social science setting, such questions also arise in statistical physics settings, where natural functions such as bond percolation turn out to be extremely sensitive to noise [BKS99]. Let us formally define noise stability.

Definition 8.1. Let $\rho \in [0, 1]$ be a correlation parameter. We say two vectors $x, y \in \{-1, 1\}^n$ are ρ-correlated if they are distributed as follows. The vector x is drawn uniformly at random, and y is obtained from x by flipping each bit x_i independently with probability $(1 - \rho)/2$. Note that for every i the correlation $E[x_i y_i] = \rho$. The *noise sensitivity* of f at ρ, $S_\rho(f)$, is simply defined as the correlation of the outputs, $E[f(x)f(y)]$.

It is not hard to see that the function maximizing noise stability is any *dictatorship* function, e.g. $f(x) = x_1$, for which $S_\rho(f) = \rho$. But another natural social scientific concern is the *influence* of players in voting schemes [BOL85], which prohibits such solutions (in democratic environments). The influence of a single voter[9] is the probability with which it can change the outcome given that all other votes are uniformly

[9]This seminal paper [BOL85] also studies the influences of coalitions of players, extremely natural in game theory, which arises

random (so, in a dictatorship it is 1 for the dictator and 0 for all others). A fair voting scheme should have no voter with high influence. As we define influence for Real-valued functions, we will use the (conditional) *variance* to measure a player's potential effect given all other (random) votes.

Definition 8.2. A function $f : \{-1,1\}^n \to \mathbb{R}$ has influence τ if for every i, $\mathrm{Var}[x_i|x_{-i}] \leq \tau$ for all i (where x_{-i} denotes the vector x without the ith coordinate).

For example, the majority function has influence $O(1/\sqrt{n})$. The question of how small the influence of a balanced function can be is extremely interesting, and leads to a highly relevant inequality for our story (both in content and techniques). As it turns out, ultimate fairness (influence $1/n$ per player) is impossible—[KKL88] show that every function has a player with nonproportional influence, at least $\Omega(\log n/n)$. At any rate, one can ask which of the functions with *small* influence is most stable, and it is natural to guess that majority should be the best[10].

The conjecture that this is the case, called the *Majority is Stablest* conjecture, arose from a completely different and surprising angle—the field of optimization, specifically "hardness of approximation". A remarkable paper [KKMO07] has shown that this conjecture implies[11] the optimality of a certain natural algorithm for approximating the *maximum cut* of a graph (i.e. the partition of vertices that maximizes the number of edges between them)[12]. This connection is highly non-trivial, but by now we have many examples showing how the analysis of certain (semidefinite programming-based) approximation algorithms for a variety of optimization problems raise many new isoperimetric questions[13], greatly enriching this field.

The Majority is Stablest conjecture was proved in a strong form by [MOO10] shortly after it was posed. Here is a formal statement (which actually works for bounded functions).

Theorem 8.3. *For every (positive correlation parameter) $\rho \geq 0$ and $\epsilon > 0$ there exists (an influence bound) $\tau = \tau(\rho, \epsilon)$ such that for every n and every $f : \{-1,1\}^n \to [-1,1]$ of influence at most τ, $S_\rho(f) \leq S_\rho(Majority_n) + \epsilon$.*

The proof reveals another angle on the story—large deviation inequalities and invariance principles. To see the connection, recall the Berry-Esseen theorem [Fel71], generalizing the standard central limit theorem to *weighted* sums of independent random signs. In this theorem, influences arise very naturally. Consider $\sum_{i=1}^n c_i x_i$. If we normalize the weights c_i to satisfy $\sum_i c_i^2 = 1$, then c_i is the influence of the ith voter, and $\tau = \max_i |c_i|$. The quality of this central limit theorem deteriorates linearly with the influence τ. Lindeberg's proof of Berry-Esseen uses an invariance principle, showing that for linear functions, the cumulative probability distribution $Pr[\sum_{i=1}^n c_i x_i \leq t]$ (for every t) is unchanged (up to τ), *regardless* of the distribution of the variables x_i, as long as they are independent and have expectation 0 and variance 1. Thus, in particular, they can be taken to be standard Gaussian, which trivializes the problem, as the weighted sum is a Gaussian as well!

To prove their theorem, [MOO10] first observed that also in the noise stability problem, the Gaussian case is simple. If the x_i, y_i are standard Gaussians with correlation ρ, the stability problem reduces to a classical result of Borell [Bor85]: that noise stability is maximized by any hyperplane through the origin. Note that here the rotational symmetry of multidimensional Gaussians, which also aids the proof, does not distinguish "dictator" functions from majority—both are such hyperplanes. Given this theorem, an invariance principle whose quality depends on τ would do the job. They next show that it is sufficient to prove the principle only for *low degree* multilinear polynomials (as the effect of noise decays with the degree). Finally, they prove this non-linear extension of Berry-Esseen for such polynomials, a form of which we state below. They also use their invariance principle to prove other conjectures, and since the publication of their paper, quite a number of further generalizations and applications were found.

in and contributes to other areas of computational complexity (including circuit complexity, learning and pseudorandomness), and raises other analytic questions which we will not discuss here.

[10] This noise sensitivity tends, as n grows, to $S_\rho(Majority_n) = \frac{2}{\pi} \arcsin \rho$.

[11] Assuming another, complexity-theoretic, conjecture called the "Unique Games" conjecture of [Kho02] (discussed already in the metric geometry section above; see also [Kho10, Wig17])

[12] Maximum Cut is a basic optimization problem whose exact complexity is $\mathcal{NP}$-complete.

[13] Many over continuous domains, like the unit cube or Gaussian space (see citeMoNe12 for one of many examples), where the connection between noise stability and isoperimetry may be ever clearer.

Theorem 8.4 [MOO10] *Let x_i be any n independent random variables with mean 0, variance 1 and bounded 3rd moments. Let g_i be n independent standard Gaussians. Let Q be any degree d multilinear n-variate polynomial of influence τ. Then for any t,*

$$|Pr[Q(x) \leq t] - Pr[Q(g) \leq t]| \leq O(d\tau^{1/d}).$$

We now only seem to be switching gears... To conclude this section, let me give one more, very different demonstration of the surprising questions (and answers) regarding noise stability and isoperimetry, arising from the very same computational considerations of optimization of hardness of approximation. Here is the question: *What is the smallest surface area of a (volume 1) body which tiles $\mathbb{R}^d$ periodically along the integer lattice $\mathbb{Z}^d$?* Namely, we seek a d-dimensional volume 1 subset $B \subseteq \mathbb{R}^d$ such that $B + \mathbb{Z}^d = \mathbb{R}^d$, such that its boundary has minimal $(d-1)$-dimensional volume[14]. Let us denote this infimum by $s(d)$. The curious reader can stop here a bit and test your intuition, what do you expect the answer to be, asymptotically in d?

Such questions originate from the late 19th century study by Thomson (later Lord Kelvin) of *foams* in 3 dimensions [Tho87], further studied, generalized and applied in mathematics, physics, chemistry, material science and even architecture. However, for this very basic question, where periodicity is defined by the simplest integer lattice, it seems that, for large d, the trivial upper and lower bounds on $s(d)$ were not improved on for over a century. The trivial upper bound on $s(d)$ is provided by the unit cube, which has surface area $2d$. The trivial lower bound on $s(d)$ comes from ignoring the tiling restriction, and considering only the volume - here the unit volume ball has the smallest surface area, $\sqrt{2\pi e d}$. Where in this quadratic range does $s(d)$ lie? In particular, can there be *"spherical cubes"*, with $s(d) = O(\sqrt{d})$?

The last question became a central issue for complexity theorists when [FKO07] related it directly to the important Unique Games conjecture, and optimal inapproximability proofs of combinatorial problems (in particular the maximum cut problem) discussed above. The nontrivial connection, which the paper elaborates and motivates, goes through attempts to find the tightest version of Raz' [Raz98] celebrated parallel repetition theorem[15]. A limit on how "strong" a parallel repetition theorem can get was again provided by Raz [Raz11]. Extending his techniques [KORW08] to the geometric setting, resolved the question above, proving that "spherical cubes" do exist!

Theorem 8.5 [KORW08] *For all d, $s(d) \leq \sqrt{4\pi d}$*

A simple proof, and various extensions of this result were given subsequently in [AK09]. We note that all known proofs are probabilistic. Giving an explicit construction that might better illustrate how a "spherical cube" (even with much worse but non-trivial surface are) looks like, seems like a challenging problem.

9 Lattice Theory

Lattices in Euclidean space are among the most "universal" objects in mathematics, in that besides being natural (e.g. arising in crystalline structures) and worthy of study in their own right, they capture a variety of problems in different fields such as number theory, analysis, approximation theory, Lie algebras, convex geometry, and more. Many of the basic results in lattice theory, as we shall see, are *existential* (namely supply no efficient means for obtaining the objects whose existence is proved), which in some cases has limited progress on these applications.

This section tells the story of one algorithm, of Lenstra, Lenstra, and Lovász [LLL82], often called the LLL algorithm, and some of its implications on these classical applications as well as modern ones in cryptography, optimization, number theory, symbolic algebra and more. But we had better define a lattice[16] first.

[14] Note that the volume of B ensures that the interiors of $B + v$ and $B + u$ are disjoint for any two distinct integer vectors $u, v \in \mathbb{Z}^d$, so this gives a tiling.

[15] A fundamental information theoretic inequality of central importance to "amplification" of Probabilistically Checkable Proofs (PCPs).

[16] We only define full-rank lattices here, which suffice for this exposition.

Let $B = \{b_1, b_2, \ldots, b_n\}$ be a basis of $\mathbb{R}^n$. Then the *lattice* $L(B)$ denotes the set (indeed, Abelian group) of all *integer* linear combinations of these vectors, i.e. $L(B) = \{\sum_i z_i b_i : z_i \in \mathbb{Z}\}$. B is also called a basis of the lattice. Naturally, a given lattice can have many different bases, e.g. the standard integer lattice in the plane, generated by $\{(0,1),(1,0)\}$, is equally well generated by $\{(999,1),(1000,1)\}$. A basic invariant associated with a lattice L is its determinant $d(L)$, which is the absolute value of $\det(B)$ for any basis B of L (this is also the volume of the fundamental parallelpiped of the lattice). For simplicity and without loss of generality, we will assume that B is normalized so that we only consider lattices L of $d(L) = 1$.

The most basic result about lattices, namely that they must contain *short* vectors (in any norm) was proved by Minkowski (who initiated Lattice Theory, and with it, the Geometry of Numbers) [Min14].

Theorem 9.1 [Min14] *Consider an arbitrary convex set K in $\mathbb{R}^n$ which is centrally symmetric*[17] *and has volume $> 2^n$. Then, every lattice L (of determinant 1) has a nonzero point in K.*

This innocent theorem, which has a simple, but *existential* (pigeonhole) proof, turns out to have numerous fundamental applications in geometry, algebra and number theory. Among famous examples this theorem yields with appropriate choice of norms and lattices, results like Dirichlet's Diophantine approximation theorem and Lagrange's four-squares theorem, and (with much more work) the finiteness of class numbers of number fields (see e.g. [PZ89]).

From now on we will focus on short vectors in the (most natural) Euclidean norm. A direct corollary of Minkowski's theorem when applying it to the cube $K = [-1,1]^n$ yields:

Corollary 9.2. *Every lattice L of determinant 1 has a nonzero point of Euclidean norm at most $\sqrt{n}$.*

Digressing a bit, we note that very recently, a century after Minkowski, a strong converse of the above corollary[18] conjectured by Dadush (see [DR16]) for *computational* motivation, has been proved in [RSD16]. This converse has many structural consequences, on the covering radius of lattices, arithmetic combinatorics, Brownian motion and others. We will not elaborate here on this new interaction of computational complexity and optimization with lattice theory and convex geometry. The papers above beautifully motivate these connections and applications, and the history of ideas and technical work needed for this complex proof.

Returning to Minkowski's corollary for the Euclidean norm, the proof is still existential, and the obvious algorithm for finding such a short vector requires exponential time in n. The breakthrough paper [LLL82] describe the LLL algorithm, an efficient, polynomial-time algorithm, which approximates the length of the shortest vector in any n-dimensional lattice by a 2^n factor.

Theorem 9.3 [LLL82] *There is a polynomial time algorithm, which given any lattice L produces a vector in L of Euclidean length at most 2^n factor longer than the shortest vector in L.*

This exponential bound may seem excessive at first, but the number and diversity of applications is staggering. First, in many problems, the dimension n is a small constant (so the actual input length arises from the bit-size of the given basis). This leads, for instance, to Lenstra's algorithm for (exactly solving) Integer Programming [Len83] in constant dimensions. It also leads to Odlyzko and Riele's refutation [OtR85] of Mertens' conjecture about cancellations in the Möbius function, and to the long list of number theoretic examples in [Sim10]. But it turns out that even when n is arbitrarily large, many problems can be solved in poly(n)-time as well. Here is a list of examples of old and new problems representing this variety, some going back to the original paper [LLL82]. In all, it suffices that real number inputs are approximated to poly(n) digits in dimension n.

- **Diophantine approximation**. While the best possible approximation of one real number by rationals with bounded denominator is readily solved by its (efficiently computable) continued fraction expansion, no such procedure is known for *simultaneous* approximation. Formally, given a *set* of real numbers, say $\{r_1, r_2, \ldots, r_n\}$, a bound Q and $\epsilon > 0$, find integers $q \leq Q$ and $p_1, \ldots, p_n$ such that all $|r_i - p_i/q| \leq \epsilon$. Existentially (using Minkowski), the Dirichlet "box-principle" shows that $\epsilon < Q^{1/n}$ is possible. Using LLL, one efficiently obtains $\epsilon < 2^{n^2} Q^{1/n}$ which is meaningful for Q described by poly(n) many bits.

[17] Namely, $x \in K$ implies that also $-x \in K$. Such sets are precisely balls of arbitrary norms.
[18] Which has to be precisely formulated.

- **Minimal polynomials of algebraic numbers**. Here we are given a single real number r and a degree bound n, and are asked if there is a polynomial $g(x)$ with integer coefficients, of degree at most n of which r is a root (and also to produce such a polynomial g if it exists). Indeed, this is a special case of the problem above with $r_i = r^i$. While the algorithm only outputs g for which $g(r) \approx 0$, it is often easy to check that it actually vanishes. Note that by varying n we can find the minimal such polynomial.

- **Polynomial factorization over Rationals**. Here the input is an integer polynomial h of degree n, and we want to factor it over $\mathbb{Q}$. The high level idea is to first find an (approximate) root r of h (e.g. using Newton's method), feed it to the problem above, which will return a minimal g having r as a root, and thus divides h. We stress that this algorithm produces the exact factorization, not an approximate one!

- **Small integer relations between reals**. Given reals $r_1, r_2, \ldots r_n$, and a bound Q, determine if there exist integers $|z_i| < Q$ such that $\sum_i z_i r_i = 0$ (and if so, find these integers). As a famous example, LLL can find an integer relation among $\arctan(1) \approx 0.785398, \arctan(1/5) \approx 0.197395$ and $\arctan(1/239) \approx 0.004184$, yielding Machin's formula

$$\arctan(1) - 4\arctan(1/5) + \arctan(1/239) = 0$$

- **Cryptanalysis**. Note that a very special case of the problem above (in which the coefficients z_i must be Boolean) is the "Knapsack problem," a famous $\mathcal{NP}$-complete problem. The point here is that in the early days of cryptography, some systems were based on the assumed "average case" hardness of Knapsack. Many such systems were broken by using LLL, e.g. [Lag84]. LLL was also used to break some versions of the RSA cryptosystem (with "small public exponents").

It is perhaps a fitting epilogue to the last item that lattices cannot only destroy cryptosystems, but also create them. The problem of efficiently approximating short vectors up to polynomial (as opposed to exponential, as LLL produces) factors is believed to be computationally hard. Here are some major consequences of this assumption. First, Ajtai showed in a remarkable paper [Ajt96] that such hardness is preserved "on average", over a cleverly-chosen distribution of random lattices. This led to a new public-key encryption scheme by Ajtai and Dwork [AD97] based on this hardness, which is arguably the only one known that can potentially sustain quantum attacks (Shor's efficient quantum algorithms can factor integers and compute discrete logarithms [Sho94]). In another breakthrough work of Gentry [Gen09], this hardness assumption is used to devise *fully homomorphic* encryption, a scheme which allows not only to encrypt data, but to perform arbitrary computations directly with encrypted data. See more in this excellent survey [Pei16].

10 Invariant Theory

Invariant theory, born in an 1845 paper of Cayley [Cay45], is major branch of algebra, with important natural connections to algebraic geometry and representation theory, but also to many other areas of mathematics. We will see some here, as well as some new connections with computational complexity, leading to new questions and results in this field. We note that computational efficiency was always important in invariant theory, which is rife with ingenious algorithms (starting with Cayley's *Omega process*), as is evident from the books [CLO92, Stu08, DK15].

Invariants are familiar enough, from examples like the following.

- In high school physics we learn that energy and momentum are preserved (namely, are *invariants*) in the dynamics of general physical systems.

- In chemical reactions the number of atoms of each element is preserved as one mixture of molecules is transformed to yield another (e.g. as combining Sodium Hydroxide ($NaOH$) and Hydrochloric Acid (HCl) yields the common salt Sodium Chloride ($NaCl$) and Water (H_2O)).

- In geometry, a classical puzzle asks when can a plane polygon be "cut and pasted" along straight lines to another polygon. Here the obvious invariant, *area*, is the only one![19]. However in generalizing this puzzle to 3-dimensional polyhedra, it turns out that besides the obvious invariant, *volume*, there is another invariant, discovered by Dehn[20].

More generally, questions about the equivalence of two surfaces (e.g. knots) under homeomorphism, whether two groups are isomorphic, or whether two points are in the same orbit of a dynamical system, etc., all give rise to similar questions and treatment. A canonical way to give negative answers to such questions is through *invariants*, namely quantities preserved under some action on an underlying space.

We will focus on invariants of *linear groups* acting on *vector spaces*. Let us present some notation. Fix a field $\mathbb{F}$ (while problems are interesting in every field, results mostly work for infinite fields only, and sometimes just for characteristic zero or algebraically closed ones). Let G be a group, and V a representation of G, namely an $\mathbb{F}$-vector space on which G acts: for every $g, h \in G$ and $v \in V$ we have $gv \in V$ and $g(hv) = (gh)v$.

The *orbit* under G of a vector (or point) $v \in V$, denoted Gv is the set of all other points hat v can be moved to by this action, namely $\{gv : g \in G\}$. Understanding the orbits of a group objects is a central task of this field. A basic question capturing many of the examples above is, given two points $u, v \in V$, do they lie in the same G-orbit, namely if $u \in Gv$. A related basic question, which is even more natural in algebraic geometry (when the field $\mathbb{F}$ is algebraically closed of characteristic zero) is whether the *closures*[21] of the two orbits intersect, namely if some point in Gu can be approximated arbitrarily well by points in Gv. We will return to specific incarnations of these questions.

When G acts on V, it also acts on $\mathbb{F}[V]$, the polynomial functions on V, also called the *coordinate ring* of V. In our setting V will have finite dimension (say m), and so $\mathbb{F}[V]$ is simply $\mathbb{F}[x_1, x_2, \ldots, x_m] = \mathbb{F}[X]$, the polynomial ring over $\mathbb{F}$ in m variables. We will denote by gp the action of a group element $g \in G$ on a polynomial $p \in \mathbb{F}[V]$.

A polynomial $p(X) \in \mathbb{F}[X]$ is *invariant* if it is unchanged by this action, namely for every $g \in G$ we have $gp = p$. All invariant polynomials clearly form a subring of $\mathbb{F}[X]$, denoted $\mathbb{F}[X]^G$, called the *ring of invariants* of this action. Understanding the invariants of group actions is the main subject of Invariant Theory. A fundamental result of Hilbert [Hil93] shows that in our linear setting[22], *all* invariant rings will be *finitely generated* as an algebra[23]. Finding the "simplest" such generating set of invariants is our main concern here.

Two familiar examples of perfect solutions to this problem follow.

- In the first, $G = S_m$, the symmetric group on m letters, is acting on the set of m formal variables X (and hence the vector space they generate) by simply permuting them. Then a set of generating invariants is simply the first m *elementary* symmetric polynomials in X.

- In the second, $G = SL_n(\mathbb{F})$, the simple linear group of matrices with determinant 1, is acting on the vector space $M_n(\mathbb{F})$ of $n \times n$ matrices (so $m = n^2$), simply by left matrix multiplication. In this case all polynomial invariants are generated by a single polynomial, the determinant of this m-variable matrix X.

In these two cases, which really supply a complete understanding of the invariant ring $\mathbb{F}[X]^G$, the generating sets are *good* in several senses. There are *few* generating invariants, they all have *low* degree, and they are *easy* to compute[24]—all these quantities are bounded by a polynomial in m, the dimension of the vector

[19] And so, every two polygons of the same area can be cut to produce an *identical* (multi)sets of triangles.

[20] So there are pairs of 3-dimensional polyhedra of the same volume, which cannot be cut to identical (multi)sets of tetrahedra.

[21] One can take closure in either the Euclidean or the Zariski topology (the equivalence in this setting proved by Mumford [Mum95]).

[22] The full generality under which this result holds is actions of *reductive* groups, which we will not define here, but includes all examples we discuss.

[23] This means that there is a finite set of polynomials $\{q_1, q_2, \ldots, q_t\}$ in $\mathbb{F}[X]^G$ so that for every polynomial $p \in \mathbb{F}[X]^G$ there is a t-variate polynomial r over $\mathbb{F}$ so that $p = r(q_1, q_2, \ldots, q_t)$.

[24] E.g. have *small* arithmetic circuits or formulae.

space[25] . In such good cases, one has efficient algorithms for the basic problems regarding orbits of group actions. For example, a fundamental duality theorem of Geometric Invariant Theory [MFK82] (see Theorem A.1.1), show how generating sets of the invariant ring can be used for the orbit closure intersection problem.

Theorem 10.1 [MFK82] *For an algebraically closed field $\mathbb{F}$ of characteristic 0, the following are equivalent for any two $u, v \in V$ and generating set P of the invariant ring $\mathbb{F}[X]^G$.*

- *The orbit closures of u and v intersect.*
- *For every polynomial $p \in P$, $p(v) = p(u)$.*

10.1 Geometric Complexity Theory (GCT)

We now briefly explain one direction from which computational complexity became interested in these algebraic problems, in work that has generated many new questions and collaboration between the fields. First, some quick background on the main problem of arithmetic complexity theory (see Chapter **??** for definitions and more discussion). In [Val79] Valiant defined arithmetic analogs $\mathcal{VP}$ and $\mathcal{VNP}$ of the complexity classes $\mathcal{P}$ and $\mathcal{NP}$ respectively, and conjectured that these two arithmetic classes are different (see Conjecture **??**). He further proved (via surprising completeness results) that to separate these classes it is sufficient to prove that the *permanent* polynomial on $n \times n$ matrices does not project to the *determinant* polynomial on $m \times m$ matrices for any $m = \text{poly}(n)$. Note that this is a pure and concrete algebraic formulation of a central computational conjecture.

In a series of papers, Mulmuley and Sohoni introduced *Geometric Complexity Theory* (GCT) to tackle this major open problem[26]. This program is surveyed by Mulmuley here [Mul12a, Mul11], as well as in Landsberg's book [Lan17]. Very concisely, the GCT program starts off as follows. First, a simple "padding" of the $n \times n$ permanent polynomial makes it have degree m and act on the entries of an $m \times m$ matrix. Consider the linear group SL_{m^2} action on all entries of such $m \times m$ matrices. This action extends to polynomials in those variables, and so in particular the two we care about: determinant and modified permanent. *The main connection is that the permanent projects to the determinant (in Valiant's sense) if and only if the orbit closures of these two polynomials intersect.* Establishing that they do not intersect (for $m = \text{poly}(n)$) naturally leads to questions about finding representation theoretic obstructions to such intersection (and hence, to the required computational lower bound). This is where things get very complicated, and describing them is beyond the scope of this survey. We note that to date, the tools of algebraic geometry and representation theory were not sufficient even to improve the quadratic bound on m of Theorem **??**. Indeed, some recent developments show severe limitations to the original GCT approach (and perhaps guiding it in more fruitful directions); see [BIP16] and its historical account. Nevertheless, this line of attack (among others in computational complexity) has lead to many new questions in computational commutative algebra and to growing collaborations between algebraists and complexity theorists – we will describe some of these now.

To do so, we will focus on two natural actions of linear groups on *tuples* of matrices, simultaneous conjugation and the left-right action. Both are special cases of *quiver representations* (see [Gab72,DW06])[27]. For both group actions we will discuss the classical questions and results on the rings of invariants, and recent advances motivated by computational considerations.

10.2 Simultaneous Conjugation

Consider the following action of $SL_n(\mathbb{F})$ on d-tuples of $n \times n$ matrices. We have $m = dn^2$ variables arranged as d $n \times n$ matrices $X = (X_1, X_2, \ldots, X_d)$. The action of a matrix $Z \in SL_n(\mathbb{F})$ on this tuple is by simultaneous

[25] There are additional desirable structural qualities of generating sets that we will not discuss, e.g. completely understanding algebraic relations between these polynomials (called *syzygies*).

[26] Origins of using invariant theory to argue computational difficulty via similar techniques go back to Strassen [Str87].

[27] We will not elaborate on the theory of quivers representation here, but only remark that reductions and completeness occur in this study as well! The left-Right quiver is *complete* in a well defined sense (see [DM15], Section 5). Informally, this means understanding its (semi)-invariants implies the same understanding of the (semi)-invariants of *all* acyclic quivers.

conjugation, by transforming it to the tuple $(Z^{-1}X_1Z, Z^{-1}X_2Z, \cdots, Z^{-1}X_dZ)$. Now, the general question above, for this action, is which polynomials in the variables X are invariant under this action?

The work of Procesi, Formanek, Razmyslov, and Donkin [Pro76, For84, Raz74, Don92] provides a good set (in most aspects discussed above) of generating invariants (over algebraically closed fields of characteristic zero). The generators are simply the traces of products of length at most n^2 of the given matrices[28]. Namely the set

$$\{Tr(X_{i_1}X_{i_2}\cdots X_{i_t}) : t \leq n^2, i_j \in [d]\}.$$

These polynomials are explicit, have small degree and are easily computable. The one shortcoming is the *exponential* size of this generating set. For example, using it to decide the intersection of orbit closures will only lead to an exponential time algorithm.

By Hilbert's existential Noether's normalization lemma [Hil93][29] we know that the size of this set of generating invariants can, in principle, be reduced to $dn^2 + 1$. Indeed, when the group action is on a vector space of dimension m, taking $m + 1$ "random" linear combinations of any finite generating set will result (with probability 1) in a small generating set. However, as we start with an exponential number of generators above, this procedure is both inefficient and also not explicit (it is not clear how to make it deterministic). One can get an explicit generating set of minimal size deterministically using the Gröbner basis algorithm (see [MR11] for the best known complexity bounds) but this will take doubly exponential time in n.

The works above [Mul12b, FS13] reduce this complexity to polynomial time! This happened in two stages. First Mulmuley [Mul12b] gave a probabilistic polynomial time algorithm, by cleverly using the structure of the exponentially many invariants above (using which one can obtain sufficiently random linear combinations using only polynomially many random bits and in polynomial time). He then argues that using conditional derandomization results, of the nature discussed in Section **??**, one can derive a deterministic polynomial time algorithm under natural computational hardness assumptions. Shortly afterwards, Forbes and Shpilka [FS13] showed that how de-randomized a variant of Mulmuley's algorithm *without* any unproven assumption, yielding an unconditional deterministic polynomial time algorithm for the problem! Their algorithm uses the derandomization methodology: very roughly speaking, they first notice that Mulmuley's probabilistic algorithm can be implemented by a very restricted computational model (a certain read-once branching program), and then use an efficient pseudo-random generator for this computational model. Here is one important algorithmic corollary (which can be extended to other quivers).

Theorem 10.2 [Mul12b, FS13] *There is a deterministic polynomial time algorithm to solve the following problem. Given two tuples of* rational *matrices* $(A_1, A_2, \ldots, A_d), (B_1, B_2, \ldots, B_d)$,, *determine if the closure of their orbits under simultaneous conjugation intersect.*

It is interesting to remark that if we only consider the orbits themselves (as opposed to their closure), namely ask if there is $Z \in SL_n(\mathbb{F})$ such that for all $i \in [d]$ we have $Z^{-1}A_iZ = B_i$, this becomes the *module isomorphism* problem over $\mathbb{F}$. For this important problem there is a deterministic algorithm (of a very different nature than above, using other algebraic tools) that can solve the problem over any field $\mathbb{F}$ using only a polynomial number of arithmetic operations over $\mathbb{F}$ [BL08].

10.3 Left-Right action

Consider now the following action of two copies, $SL_n(\mathbb{F}) \times SL_n(\mathbb{F})$ on d-tuples of $n \times n$ matrices. We still have $m = dn^2$ variables arranged as d $n \times n$ matrices $X = (X_1, X_2, \ldots, X_d)$. The action of a pair of matrices $Z, W \in SL_n(\mathbb{F})$ on this tuple is by left-right action, transforming it to the tuple $(Z^{-1}X_1W, Z^{-1}X_2W, \cdots, Z^{-1}X_dW)$. Again, for this action, is which polynomials in the variables X are invariant under this action? Despite the superficial similarity to the to simultaneous conjugation, the invariants here have entirely different structure, and bounding their size required different arguments.

[28] Convince yourself that such polynomials are indeed invariant.

[29] We remark that this is the same foundational paper which proved the *finite basis* and *Nullstellensatz* theorems. It is interesting that Hilbert's initial motivation to formulate and prove these cornerstones of commutative algebra was the search for invariants of linear actions.

The works of [DW00, DZ01, SVdB01, ANS10] provides an infinite set of generating invariants. The generators (again, over algebraically closed fields) are determinants of linear forms of the d matrices, with *matrix* coefficients of arbitrary dimension. Namely the set

$$\{\det(C_1 \otimes X_1 + C_2 \otimes X_2 + \cdots + C_d \otimes X_d) \,:\, C_i \in M_k(\mathbb{F}), k \in \mathbb{N}\}.$$

These generators, while concisely described, fall short on most goodness aspects above, and we now discuss improvements. First, by Hilbert's finite generation, we know in particular that some finite bound k on the dimension of the matrix coefficients C_i exist. A quadratic upper bound $k \leq n^2$ was obtained by Derksen and Makam [DM15] after a long sequence of improvements described there. Still, there is an exponential number[30] of possible matrix coefficients of this size which can be described explicitly, and allowing randomness one can further reduce this number to a polynomial. Thus we e.g. have the following weaker analog to the theorem above regarding orbit closure intersection for this left-right action.

Theorem 10.3. *There is a probabilistic polynomial time algorithm to solve the following problem. Given two tuples of* rational *matrices* $(A_1, A_2, \ldots, A_d), (B_1, B_2, \ldots, B_d)$*, determine if the closure of their orbits under the left-right action intersect.*

In the remainder we discuss an important special case of this problem, namely when all $B_i = 0$, for which a *deterministic* polynomial time algorithm was found. While this problem is in commutative algebra, this algorithm surprisingly has implications in analysis and non-commutative algebra, and beyond to computational complexity and quantum information theory. We will mention some of these, but let us start by defining the problem.

For an action of a linear group G on a vector space V, define the *nullcone* of the action to be the set of all points $v \in V$ such that the the closure of the orbit Gv contains 0. The points in the nullcone are sometimes called *unstable*. The nullcone of fundamental importance in invariant theory! Some examples of nullcones for actions we have discussed are the following. For the action of $SL_n(\mathbb{C})$ on $M_n(\mathbb{C})$ by left multiplication, it is the set of *singular* matrices. For the action of $SL_n(\mathbb{C})$ on $M_n(\mathbb{C})$ by conjugation, it is the set of *nilpotent* matrices. As you would guess (one direction is trivial), the nullcone is precisely the set of points which vanish under all invariant polynomials. Thus if we have a good generating set one can use them to efficiently test membership in the nullcone. However, we are not in this situation for the left-right action. Despite that a derterministic polynomial time algorithm was obtained in [GGOW15] over the complex numbers, and then a very different algorithm by [IQS15] which works for all fields. These two algorithms have different nature and properties, and use in different ways the upper bounds on the dimension of matrix coefficients in the invariants.

Theorem 10.4 [GGOW15, IQS15] *There is a deterministic polynomial time algorithm, that on a given a tuple of matrices* $(A_1, A_2, \ldots, A_d)$ *in* $M_n(\mathbb{F})$ *determines if it is in the nullcone of the left-right action.*

We conclude with some of the diverse consequences of this algorithm. All the precise definitions of the notions below, as well as the proofs, interconnections and the meandering story leading to these algorithms can be found in [GGOW15, GGOW16].

Theorem 10.5 [GGOW15, GGOW16] *There are deterministic polynomial time algorithms to solve the following problems.*

- *The feasibility problem for Brascamp-Lieb ineuqalities, and more generally, computing the optimal constant for each.*
- *The word problem over the (non-commutative) free skew field.*
- *Computing the non-commutative rank of a symbolic matrix*[31]*.*
- *Approximating the commutative rank of a symbolic matrix to within a factor of two.*[32]*.*
- *Testing if a completely positive quantum operator is rank-decreasing.*

[30] Well, a possible infinite number, but it can be reduced to exponential.
[31] A matrix whose entries are linear forms in a set of variables
[32] Computing this rank exactly is the PIT problem discussed at the end of Section **??**.

References

[AB03] M. Agrawal and S. Biswas. Primality and identity testing via Chinese remaindering. *Journal of the ACM*, 50(4):429–443, 2003. 3

[AB09] S. Arora and B. Barak. *Computational complexity: a modern approach*. Cambridge University Press, 2009. 9

[ACORT11] D. Achlioptas, A. Coja-Oghlan, and F. Ricci-Tersenghi. On the solution-space geometry of random constraint satisfaction problems. *Random Structures & Algorithms*, 38(3):251–268, 2011. 11

[AD97] M. Ajtai and C. Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pages 284–293. ACM, 1997. 15

[AH92] L. M. Adleman and M. A. Huang. *Primality testing and abelian varieties over finite fields*, volume 1512. Springer, 1992. 3

[Ajt96] M. Ajtai. Generating hard instances of lattice problems. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 99–108. ACM, 1996. 15

[AK09] Noga Alon and Bo'az Klartag. Economical toric spines via cheeger's inequality. *Journal of Topology and Analysis*, 1(02):101–111, 2009. 13

[AKS04] M. Agrawal, N. Kayal, and N. Saxena. Primes is in *P*. *Ann. of Math.*, 160(2):781–793, 2004. 3

[Ald83] D. J. Aldous. Random walks on finite groups and rapidly mixing Markov chains. In *Séminaire de Probabilités XVII 1981/82*, pages 243–297. Springer, 1983. 10

[Ald90] D. J. Aldous. The random walk construction of uniform spanning trees and uniform labelled trees. *SIAM Journal on Discrete Mathematics*, 3(4):450–465, 1990. 10

[ALN08] Sanjeev Arora, James Lee, and Assaf Naor. Euclidean distortion and the sparsest cut. *Journal of the American Mathematical Society*, 21(1):1–21, 2008. 7

[ANS10] Bharat Adsul, Suresh Nayak, and KV Subrahmanyam. A geometric approach to the kronecker problem ii: Invariants of matrices for simultaneous left-right actions. *Manuscript, available in http://www. cmi. ac. in/kv/ANS10. pdf*, 2010. 19

[APR83] L. Adleman, C. Pomerance, and R. S. Rumely. On distinguishing prime numbers from composite numbers. *Annals of Mathematics*, pages 173–206, 1983. 3

[Bab91] L. Babai. Local expansion of vertex-transitive graphs and random generation in finite groups. In *STOC*, volume 91, pages 164–174. Citeseer, 1991. 9

[Bab15] László Babai. Graph isomorphism in quasipolynomial time. *arXiv preprint arXiv:1512.03547*, 2015. 8

[BDWY13] B. Barak, Z. Dvir, A. Wigderson, and A. Yehudayoff. Fractional Sylvester–Gallai theorems. *Proceedings of the National Academy of Sciences*, 110(48):19213–19219, 2013. 5

[Ber67] E. R. Berlekamp. Factoring polynomials over finite fields. *Bell System Technical Journal*, 46(8):1853–1859, 1967. 4

[Bes19] A. Besicovitch. Sur deux questions d'integrabilité des fonctions. *J. Soc. Phys. Math.*, 2:105–123, 1919. 4

[BIP16] Peter Bürgisser, Christian Ikenmeyer, and Greta Panova. No occurrence obstructions in geometric complexity theory. *arXiv preprint arXiv:1604.06431*, 2016. 17

[BIW06] B. Barak, R. Impagliazzo, and A. Wigderson. Extracting randomness using few independent sources. *SIAM Journal on Computing*, 36(4):1095–1118, 2006. 5

[BKS99] I. Benjamini, G. Kalai, and O. Schramm. Noise sensitivity of Boolean functions and applications to percolation. *Publications Mathématiques de l'Institut des Hautes Études Scientifiques*, 90(1):5–43, 1999. 11

[BKT04] J. Bourgain, N. Katz, and T. Tao. A sum-product estimate in finite fields, and applications. *Geometric & Functional Analysis GAFA*, 14(1):27–57, 2004. 5

[BL06] Yonatan Bilu and Nathan Linial. Lifts, discrepancy and nearly optimal spectral gap. *Combinatorica*, 26(5):495–519, 2006. 6

[BL08] Peter A Brooksbank and Eugene M Luks. Testing isomorphism of modules. *Journal of Algebra*, 320(11):4020–4029, 2008. 18

[BLG12] Henrik Bäärnhielm and Charles R Leedham-Green. The product replacement prospector. *Journal of Symbolic Computation*, 47(1):64–75, 2012. 9

[BOL85] M. Ben-Or and N. Linial. Collective coin flipping, robust voting schemes and minima of Banzhaf values. In *Foundations of Computer Science, 1985., 26th Annual Symposium on*, pages 408–416. IEEE, 1985. 11

[Bor85] C. Borell. Geometric bounds on the Ornstein-Uhlenbeck velocity process. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 70(1):1–13, 1985. 12

[Bou85] J. Bourgain. On Lipschitz embedding of finite metric spaces in Hilbert space. *Israel Journal of Mathematics*, 52(1-2):46–52, 1985. 6

[Bro89] A. Broder. Generating random spanning trees. In *Foundations of Computer Science, 1989., 30th Annual Symposium on*, pages 442–447. IEEE, 1989. 10

[BS84] L. Babai and E. Szemerédi. On the complexity of matrix group problems I. In *Foundations of Computer Science, 1984. 25th Annual Symposium on*, pages 229–240. IEEE, 1984. 9

[BS97] E. Bach and J. Shallit. Algorithmic number theory. Efficient algorithms, vol. 1, 1997. 4

[BSS14] J. Batson, D. A. Spielman, and N. Srivastava. Twice-Ramanujan sparsifiers. *SIAM Review*, 56(2):315–334, 2014. 6

[BT91] J. Bourgain and L. Tzafriri. On a problem of Kadison and Singer. *Journal Fur Die Reine Und Angewandte Mathematik*, 420:1–43, 1991. 5

[Cay45] Arthur Cayley. *On the theory of linear transformations*. Number 1. E. Johnson, 1845. 15

[CLO92] D. Cox, J. Little, and D. O'Shea. Ideals, varieties, and algorithms: an introduction to computational algebraic geometry and commutative algebra, 1992. 15

[Coo02] G. Cooperman. Towards a practical, theoretically sound algorithm for random generation in finite groups. *arXiv preprint math/0205203*, 2002. 9

[Dav71] R. O. Davies. Some remarks on the Kakeya problem. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 69, pages 417–421. Cambridge Univ Press, 1971. 4

[DFK91] M. Dyer, A. Frieze, and R. Kannan. A random polynomial-time algorithm for approximating the volume of convex bodies. *Journal of the ACM (JACM)*, 38(1):1–17, 1991. 10

[Dia88] P. Diaconis. Group representations in probability and statistics. *Lecture Notes-Monograph Series*, pages i–192, 1988. 10

[Dix08] J. D. Dixon. Generating random elements in finite groups. *The electronic journal of combinatorics*, 13(R94):1, 2008. 9

[DK15] Harm Derksen and Gregor Kemper. *Computational invariant theory*. Springer, 2015. 15

[DKSS13] Z. Dvir, S. Kopparty, S. Saraf, and M. Sudan. Extensions to the method of multiplicities, with applications to Kakeya sets and mergers. *SIAM Journal on Computing*, 42(6):2305–2328, 2013. 5

[DM15] Harm Derksen and Visu Makam. Polynomial degree bounds for matrix semi-invariants. *arXiv preprint arXiv:1512.03393*, 2015. 17, 19

[Don92] S. Donkin. Invariants of several matrices. *Inventiones mathematicae*, 110(1):389–401, 1992. 18

[DR16] Daniel Dadush and Oded Regev. Towards strong reverse Minkowski-type inequalities for lattices. In *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*, pages 447–456. IEEE, 2016. 14

[DS07] Z. Dvir and A. Shpilka. Locally decodable codes with two queries and polynomial identity testing for depth 3 circuits. *SIAM Journal on Computing*, 36(5):1404–1434, 2007. 5

[DSVW04] Martin Dyer, Alistair Sinclair, Eric Vigoda, and Dror Weitz. Mixing in time and space for lattice spin systems: A combinatorial view. *Random Structures & Algorithms*, 24(4):461–479, 2004. 11

[Dvi09] Z. Dvir. On the size of Kakeya sets in finite fields. *Journal of the American Mathematical Society*, 22(4):1093–1097, 2009. 5

[Dvi10] Z. Dvir. Guest column: from randomness extraction to rotating needles. *ACM SIGACT News*, 40(4):46–61, 2010. 5

[DW00] Harm Derksen and Jerzy Weyman. Semi-invariants of quivers and saturation for Littlewood-Richardson coefficients. *Journal of the American Mathematical Society*, 13(3):467–479, 2000. 19

[DW06] Harm Derksen and Jerzy Weyman. The combinatorics of quiver representations. *arXiv preprint math/0608288*, 2006. 17

[DZ01] Matyas Domokos and Alexander N Zubkov. Semi-invariants of quivers as determinants. *Transformation groups*, 6(1):9–24, 2001. 19

[Fel71] W. Feller. An introduction to probability theory and its applications. *Wiley series in probability and mathematical statistics.*, 1971. 12

[FKO07] Uriel Feige, Guy Kindler, and Ryan O'Donnell. Understanding parallel repetition requires understanding foams. In *Computational Complexity, 2007. CCC'07. Twenty-Second Annual IEEE Conference on*, pages 179–192. IEEE, 2007. 13

[For84] E. Formanek. Invariants and the ring of generic matrices. *Journal of Algebra*, 89(1):178–223, 1984. 18

[FS13] M. A. Forbes and A. Shpilka. Explicit noether normalization for simultaneous conjugation via polynomial identity testing. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 527–542. Springer, 2013. 18

[Gab72] P. Gabriel. Unzerlegbare darstellungen I. *Manuscripta Mathematica*, 6(1):71–103, 1972. 17

[Gen09] C. Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, volume 9, pages 169–178, 2009. 15

[GGOW15] Ankit Garg, Leonid Gurvits, Rafael Oliveira, and Avi Wigderson. A deterministic polynomial time algorithm for non-commutative rational identity testing. *arXiv preprint arXiv:1511.03730*, 2015. 19

[GGOW16] Ankit Garg, Leonid Gurvits, Rafael Oliveira, and Avi Wigderson. Algorithmic aspects of brascamp-lieb inequalities. *arXiv preprint arXiv:1607.06711*, 2016. 19

[GJL16] Heng Guo, Mark Jerrum, and Jingcheng Liu. Uniform sampling through the Lovász local lemma. *arXiv preprint arXiv:1611.01647*, 2016. 11

[GK86] S. Goldwasser and J. Kilian. Almost all primes can be quickly certified. In *Proceedings of the eighteenth annual ACM symposium on Theory of computing*, pages 316–329. ACM, 1986. 3

[GK10] L. Guth and N. H. Katz. On the Erdős distinct distance problem in the plane. *arXiv preprint arXiv:1011.4105*, 2010. 5

[GLR10] V. Guruswami, J. R. Lee, and A. Razborov. Almost Euclidean subspaces of ℓ_1^n via expander codes. *Combinatorica*, 30(1):47–68, 2010. 5

[GMR89] S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proof systems. *SIAM J. Comput.*, 18(1):186–208, 1989. 9

[Goe97] M. X. Goemans. Semidefinite programming in combinatorial optimization. *Mathematical Programming*, 79(1-3):143–161, 1997. 7

[Gol08] O. Goldreich. *Computational Complexity: A Conceptual Perspective*. Cambridge University Press, Cambridge, 2008. 9

[Gra05] A. Granville. It is easy to determine whether a given integer is prime. *Bull. Amer. Math. Soc.*, 42:3–38, 2005. 3

[Gro87] M. Gromov. Essays in group theory. *Math. Sci. Res. Inst. Publ*, 8:75–263, 1987. 7

[Gur08] Leonid Gurvits. Van der Waerden/Schrijver-Valiant like conjectures and stable (aka hyperbolic) homogeneous polynomials: one theorem for all. *the electronic journal of combinatorics*, 15(1):R66, 2008. 6

[HEO05] D. F. Holt, B. Eick, and E. A. O'Brien. *Handbook of computational group theory*. CRC Press, 2005. 8

[Hil93] D. Hilbert. Über die vollen Invariantensysteme. *Mathematische Annalen*, 42(3):313–373, 1893. 16, 18

[HL72] O. J. Heilmann and E. H. Lieb. Theory of monomer-dimer systems. *Communications in Mathematical Physics*, 25(3):190–232, 1972. 9

[HLW06] S. Hoory, N. Linial, and A. Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006. 6

[IQS15] Gábor Ivanyos, Youming Qiao, and KV Subrahmanyam. Non-commutative Edmonds' problem and matrix semi-invariants. *arXiv preprint arXiv:1508.00690*, 2015. 19

[IW97] R. Impagliazzo and A. Wigderson. P = BPP unless E has subexponential circuits: Derandomizing the XOR lemma. *Proceedings of the 29th annual ACM Symposium on Theory of Computing, ACM Press, New York*, pages 220–229, 1997. 3

[JS93] Mark Jerrum and Gregory B Sorkin. Simulated annealing for graph bisection. In *Foundations of Computer Science, 1993. Proceedings., 34th Annual Symposium on*, pages 94–103. IEEE, 1993. 11

[JS96] M. Jerrum and A. Sinclair. The Markov chain Monte Carlo method: an approach to approximate counting and integration. *Approximation algorithms for NP-hard problems*, pages 482–520, 1996. 10

[JSV04] M. Jerrum, A. Sinclair, and E. Vigoda. A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries. *Journal of the ACM (JACM)*, 51(4):671–697, 2004. 10

[JVV86] M. R. Jerrum, L. G. Valiant, and V. V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoretical Computer Science*, 43:169–188, 1986. 10

[Kho02] S. Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*, pages 767–775. ACM, 2002. 7, 12

[Kho10] S. Khot. Inapproximability of NP-complete problems, discrete Fourier analysis, and geometry. In *International Congress of Mathematics*, volume 5, 2010. 7, 12

[KKL88] J. Kahn, G. Kalai, and N. Linial. The influence of variables on Boolean functions. In *Foundations of Computer Science, 1988., 29th Annual Symposium on*, pages 68–80. IEEE, 1988. 12

[KKMO07] S. Khot, G. Kindler, E. Mossel, and R. O'Donnell. Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? *SIAM Journal on Computing*, 37(1):319–357, 2007. 12

[KLMN05] Robert Krauthgamer, James R Lee, Manor Mendel, and Assaf Naor. Measured descent: A new embedding method for finite metrics. *Geometric & Functional Analysis GAFA*, 15(4):839–858, 2005. 7

[KORW08] Guy Kindler, Ryan O'Donnell, Anup Rao, and Avi Wigderson. Spherical cubes and rounding in high dimensions. In *Foundations of Computer Science, 2008. FOCS'08. IEEE 49th Annual IEEE Symposium on*, pages 189–198. IEEE, 2008. 13

[KS59] R. V. Kadison and I. M. Singer. Extensions of pure states. *American journal of mathematics*, pages 383–400, 1959. 5

[KS09] Neeraj Kayal and Shubhangi Saraf. Blackbox polynomial identity testing for depth 3 circuits. In *Foundations of Computer Science, 2009. FOCS'09. 50th Annual IEEE Symposium on*, pages 198–207. IEEE, 2009. 5

[KV05] S. A. Khot and N. K. Vishnoi. The unique games conjecture, integrality gap for cut problems and embeddability of negative type metrics into ℓ_1. In *46th Annual IEEE Symposium on Foundations of Computer Science, 2005.*, pages 53–62. IEEE, 2005. 7

[KY06] G. Kasparov and G. Yu. The coarse geometric Novikov conjecture and uniform convexity. *Advances in Mathematics*, 206(1):1–56, 2006. 7

[Laf08] V. Lafforgue. Un renforcement de la propriété (T). *Duke Mathematical Journal*, 143(3):559–602, 2008. 8

[Lag84] J. C. Lagarias. Knapsack public key cryptosystems and Diophantine approximation. In *Advances in cryptology*, pages 3–23. Springer, 1984. 15

[Lan17] Joseph Landsberg. *Geometry and complexity theory*. Cambridge University press, 2017. 17

[Len83] Jr. H. W. Lenstra. Integer programming with a fixed number of variables. *Mathematics of operations research*, 8(4):538–548, 1983. 14

[Lin02] N. Linial. Finite metric spaces-combinatorics, geometry and algorithms. In *In Proceedings of the International Congress of Mathematicians III*, pages 573–586. Citeseer, 2002. 7

[LLL82] A. K. Lenstra, H. W. Lenstra, and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, 1982. 13, 14

[LLR95] Nathan Linial, Eran London, and Yuri Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995. 6

[LPS88] A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988. 6

[LRVW03] C. Lu, O. Reingold, S. Vadhan, and A. Wigderson. Extractors: Optimal up to constant factors. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 602–611. ACM, 2003. 4, 5

[Mar88] G. A. Margulis. Explicit group-theoretic constructions of combinatorial schemes and their applications in the construction of expanders and concentrators. *Problems Information Transmission*, 24:39–46, 1988. 6

[MFK82] David Mumford, John Fogarty, and Frances Kirwan. Geometric invariant theory, 1982. 17

[Mil76] G. L. Miller. Riemann's hypothesis and tests for primality. *Journal of computer and system sciences*, 13(3):300–317, 1976. 3

[Min10] H. Minkowski. *Geometrie der Zahlen*. Teubner, 1910. 14

[MN14] M. Mendel and A. Naor. Nonlinear spectral calculus and super-expanders. *Publications mathématiques de l'IHÉS*, 119(1):1–95, 2014. 8

[Moi16] Ankur Moitra. Approximate counting, the Lovász local lemma and inference in graphical models. *arXiv preprint arXiv:1610.04317*, 2016. 11

[MOO10] E. Mossel, R. O'Donnell, and K. Oleszkiewicz. Noise stability of functions with low influences: Invariance and optimality. *Annals of mathematics*, 171(1):295–341, 2010. 12, 13

[Mos09] Robin A Moser. A constructive proof of the Lovász local lemma. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 343–350. ACM, 2009. 11

[MPZ02] M. Mézard, G. Parisi, and R. Zecchina. Analytic and algorithmic solution of random satisfiability problems. *Science*, 297(5582):812–815, 2002. 11

[MR11] E. W. Mayr and S. Ritscher. Space-efficient Gröbner basis computation without degree bounds. In *Proceedings of the 36th international symposium on Symbolic and algebraic computation*, pages 257–264. ACM, 2011. 18

[MSS13a] A. Marcus, D. A. Spielman, and N. Srivastava. Interlacing families I: Bipartite Ramanujan graphs of all degrees. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 529–537. IEEE, 2013. 6

[MSS13b] A. Marcus, D. A. Spielman, and N. Srivastava. Interlacing families II: Mixed characteristic polynomials and the Kadison-Singer problem. *arXiv preprint arXiv:1306.3969*, 2013. 5, 6

[MT10] Robin A Moser and Gábor Tardos. A constructive proof of the general Lovász local lemma. *Journal of the ACM (JACM)*, 57(2):11, 2010. 11

[Mul11] K. D. Mulmuley. On P vs. NP and geometric complexity theory. *Journal of the ACM (JACM)*, 58(2):5, 2011. 17

[Mul12a] K. D. Mulmuley. The GCT program toward the P vs. NP problem. *Communications of the ACM*, 55(6):98–107, 2012. 17

[Mul12b] K. D. Mulmuley. Geometric complexity theory V: Equivalence between blackbox derandomization of polynomial identity testing and derandomization of Noether's normalization lemma. In *Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on*, pages 629–638. IEEE, 2012. 18

[Mum95] David Mumford. *Algebraic Geometry: Complex projective varieties*, volume 1. Springer Science & Business Media, 1995. 16

[NY17] Assaf Naor and Robert Young. Vertical perimeter versus horizontal perimeter. *arXiv preprint arXiv:1701.00620*, 2017. 7

[O'D14] R. O'Donnell. *Analysis of Boolean functions*. Cambridge University Press, 2014. 11

[OtR85] A. M. Odlyzko and H. JJ. te Riele. Disproof of the Mertens conjecture. *J. reine angew. Math*, 357:138–160, 1985. 14

[Pei16] Chris Peikert. A decade of lattice cryptography. *Foundations and Trends in Theoretical Computer Science*, 10(4):283–424, 2016. 15

[Pro76] C. Procesi. The invariant theory of $n \times n$ matrices. *Advances in Mathematics*, 19(3):306–381, 1976. 18

[PZ89] M. Pohst and H. Zassenhaus. Algorithmic algebraic number theory. *Cambridge University Press*, 1989. 14

[Rab80] M. O. Rabin. Probabilistic algorithm for testing primality. *Journal of number theory*, 12(1):128–138, 1980. 3

[Raz74] Ju. P. Razmyslov. Trace identities of full matrix algebras over a field of characteristic zero. *Mathematics of the USSR-Izvestiya*, 8(4):727, 1974. 18

[Raz98] Ran Raz. A parallel repetition theorem. *SIAM Journal on Computing*, 27(3):763–803, 1998. 13

[Raz11] Ran Raz. A counterexample to strong parallel repetition. *SIAM Journal on Computing*, 40(3):771–777, 2011. 13

[RSD16] Oded Regev and Noah Stephens-Davidowitz. A reverse Minkowski theorem. *arXiv preprint arXiv:1611.05979*, 2016. 14

[RVW02] O. Reingold, S. Vadhan, and A. Wigderson. Entropy waves, the zig-zag graph product, and new constant-degree expanders. *Annals of Mathematics*, pages 157–187, 2002. 8

[Ser03] Ákos Seress. *Permutation group algorithms*, volume 152. Cambridge University Press, 2003. 8

[Sho94] P. W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *Foundations of Computer Science, 1994 Proceedings., 35th Annual Symposium on*, pages 124–134. IEEE, 1994. 4, 15

[Sim10] D. Simon. Selected applications of LLL in number theory. In *The LLL Algorithm*, pages 265–282. Springer, 2010. 14

[SJ89] Alistair Sinclair and Mark Jerrum. Approximate counting, uniform generation and rapidly mixing markov chains. *Information and Computation*, 82(1):93–133, 1989. 10

[Sly10] Allan Sly. Computational transition at the uniqueness threshold. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 287–296. IEEE, 2010. 11

[SS77] R. M. Solovay and V. Strassen. A fast Monte-Carlo test for primality. *SIAM J. Comput.*, 6(1):84–85, 1977. 3

[SS12] D. A. Spielman and N. Srivastava. An elementary proof of the restricted invertibility theorem. *Israel Journal of Mathematics*, 190(1):83–91, 2012. 6

[SSU04] Arora S., Rao S., and Vazirani U. Expander rows, geometric embeddings and graph partitioning. *Proc. of the 36th Annual Symposium on Theory of Computing*, pages 222–231, 2004. 7

[ST11] D. A. Spielman and S. Teng. Spectral sparsification of graphs. *SIAM Journal on Computing*, 40(4):981–1025, 2011. 5

[STJ83] E. Szemerédi and W. T. Trotter Jr. Extremal problems in discrete geometry. *Combinatorica*, 3(3-4):381–392, 1983. 5

[Str87] Volker Strassen. Relative bilinear complexity and matrix multiplication. *Journal für die reine und angewandte Mathematik*, 375:406–443, 1987. 17

[Stu08] Bernd Sturmfels. *Algorithms in invariant theory*. Springer Science & Business Media, 2008. 15

[SVdB01] Aidan Schofield and Michel Van den Bergh. Semi-invariants of quivers for arbitrary dimension vectors. *Indagationes Mathematicae*, 12(1):125–138, 2001. 19

[SY10] A. Shpilka and A. Yehudayoff. Arithmetic circuits: A survey of recent results and open questions. *Foundations and Trends in Theoretical Computer Science*, 5(3–4):207–388, 2010. 10

[Tao09] T. Tao. Recent progress on the Kakeya conjecture, 2009. http://terrytao.wordpress.com/2009/05/11/recent-progress-on-the-kakeya-conjecture. 4

[Tho87] William Sir Thomson. On the division of space with minimum partitional area. *Acta mathematica*, 11(1-4):121–134, 1887. 13

[Vad11] S. P. Vadhan. Pseudorandomness. *Foundations and Trends in Theoretical Computer Science*, 7(1–3):1–336, 2011. 5

[Val79] Leslie G. Valiant. The complexity of computing the permanent. *Theoretical computer science*, 8(2):189–201, 1979. 10, 17

[Vem05] Santosh Vempala. Geometric random walks: a survey. *Combinatorial and computational geometry*, 52(573-612):2, 2005. 10

[Wei06] Dror Weitz. Counting independent sets up to the tree threshold. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 140–149. ACM, 2006. 11

[Wig17] A. Wigderson. *Mathematics and Computation*. Princeton University Press, 2017. To appear. Draft available here: https://www.math.ias.edu/avi/book. 2, 3, 4, 6, 7, 9, 10, 12

[Wol99] T. Wolff. Recent work connected with the Kakeya problem. *Prospects in mathematics (Princeton, NJ, 1996)*, 2:129–162, 1999. 4