

Generalized Kakeya sets for polynomial evaluation and faster computation of fermionants

Andreas Björklund¹, Petteri Kaski², and Ryan Williams³

1 Department of Computer Science, Lund University, Lund, Sweden

2 Department of Computer Science, Aalto University, Helsinki, Finland

3 Department of Electrical Engineering and Computer Science & CSAIL, MIT, Cambridge, MA, USA

Abstract

We present two new data structures for computing values of an n -variate polynomial P of degree at most d over a finite field of q elements. Assuming that d divides $q - 1$, our first data structure relies on $(d + 1)^{n+2}$ tabulated values of P to produce the value of P at any of the q^n points using $O(nqd^2)$ arithmetic operations in the finite field. Assuming that s divides d and d/s divides $q - 1$, our second data structure assumes that P satisfies a degree-separability condition and relies on $(d/s + 1)^{n+s}$ tabulated values to produce the value of P at any point using $O(nq^s sq)$ arithmetic operations. Our data structures are based on generalizing upper-bound constructions due to Mockenhaupt and Tao (2004), Saraf and Sudan (2008), and Dvir (2009) for Kakeya sets in finite vector spaces from linear to higher-degree polynomial curves.

As an application we show that the new data structures enable a faster algorithm for computing integer-valued *fermionants*, a family of self-reducible polynomial functions introduced by Chandrasekharan and Wiese (2011) that captures numerous fundamental algebraic and combinatorial invariants such as the determinant, the permanent, the number of Hamiltonian cycles in a directed multigraph, as well as certain partition functions of strongly correlated electron systems in statistical physics. In particular, a corollary of our main theorem for fermionants is that the permanent of an $m \times m$ integer matrix with entries bounded in absolute value by a constant can be computed in time $2^{m-\Omega(\sqrt{m/\log \log m})}$, improving an earlier algorithm of Björklund (2016) that runs in time $2^{m-\Omega(\sqrt{m/\log m})}$.

1998 ACM Subject Classification F.2.1, F.2.2, G.2.1, G.2.2

Keywords and phrases Besicovitch set, fermionant, finite field, finite vector space, Hamiltonian cycle, homogeneous polynomial, Kakeya set, permanent, polynomial evaluation, tabulation

Digital Object Identifier 10.4230/LIPIcs.IPEC.2017.7

1 Introduction

The protagonist of this paper is the following task. We want an efficient representation of an n -variate degree- d polynomial P over a finite field \mathbb{F}_q of order q , that permits us to evaluate P on arbitrary points $a \in \mathbb{F}_q^n$. What kind of resource trade-offs can be achieved between space (for representing P) and query time (for computing $P(a)$ at a given a)?

The study of data structures that enable fast “polynomial evaluation” queries for multivariate polynomials was initiated by Kedlaya and Umans [11] for polynomials with bounded individual variable degrees, motivated by applications to fast polynomial factorization. (For *univariate* polynomial evaluation, cf. von zur Gathen and Gerhard [24].) Here we focus on



© Andreas Björklund, Petteri Kaski, Ryan Williams;
licensed under Creative Commons License CC-BY

12th International Symposium on Parameterized and Exact Computation (IPEC 2017).

Editors: Daniel Lokshantov and Naomi Nishimura; Article No. 7; pp. 7:1–7:12

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

the case when P has (total) degree d , in particular, when d is less than n .¹

We seek data structures consisting of a set $K \subseteq \mathbb{F}_q^n$ and an associated list $((a, P(a)) : a \in K)$ of evaluations. There are two extremes for such designs. At one extreme, we can set $K = \mathbb{F}_q^n$, put all evaluations in a sorted array, and binary search achieves $O(n \log q)$ query time. At the other extreme, to uniquely identify P we must tabulate $\Omega(\binom{n+d}{d})$ points, as this is the dimension of the monomial basis. However, when K is this small, we are only aware of brute-force ($n^{O(d)}$ -time) algorithms to evaluate the polynomial in any other point. Between these two extremes, we seek constructions for sets K that suffice for evaluating P at any point outside K in time that scales sub-exponentially in d . Our motivation is to accelerate the best known algorithms for canonical #P-hard problems (cf. Section 1.2).

1.1 Polynomial evaluation based on generalized Kakeya sets

Let $\mathbb{F}_q[x]$ be the ring of polynomials over indeterminates $x = (x_1, x_2, \dots, x_n)$ with coefficients in \mathbb{F}_q . Our first main theorem constructs an explicit set $K \subseteq \mathbb{F}_q^n$ of cardinality at most $(d+1)^{n+2}$ which allows for relatively quick evaluation of any degree- d P at all points in \mathbb{F}_q^n .

► **Theorem 1.** *Let d divide $q-1$. There is a set $K \subseteq \mathbb{F}_q^n$ of size $|K| \leq (d+1)^{n+2}$ along with functions $g_1, g_2, \dots, g_{(q-1)(d+1)^2} : \mathbb{F}_q^n \rightarrow K$ and scalars $\gamma_1, \gamma_2, \dots, \gamma_{(q-1)(d+1)^2} \in \mathbb{F}_q$ such that for every polynomial $P \in \mathbb{F}_q[x]$ of degree at most d and every vector $a \in \mathbb{F}_q^n$,*

$$P(a) = \sum_{j=1}^{(q-1)(d+1)^2} \gamma_j P(g_j(a)).$$

► **Remark.** Let us write $M(q)$ for the time complexity² of multiplication and division in \mathbb{F}_q . The construction in Theorem 1 is *explicit* in the sense that (a) there is an algorithm that in time $O(|K|nqM(q))$ lists the elements of K ; and (b) there is an algorithm that in time $O(nqd^2M(q))$ computes the values $g_j(a) \in \mathbb{F}_q^n$ and $\gamma_j \in \mathbb{F}_q$ for all $j = 1, 2, \dots, (q-1)(d+1)^2$ when given $a \in \mathbb{F}_q^n$ as input. The quadratic dependence on d has not been optimized.

The size of K can be further reduced for polynomials P satisfying a certain (natural) restriction which holds for several well-studied polynomials. Suppose we partition the variable set $X = \{x_1, x_2, \dots, x_n\}$ into $X = X_1 \cup X_2 \cup \dots \cup X_d$ such that $|X_1| = |X_2| = \dots = |X_d| = n/d$. Let us say that a degree- d polynomial $P \in \mathbb{F}_q[x]$ is *degree-separable* relative to X_1, X_2, \dots, X_d if every monomial of P contains one variable from each X_i . Note a degree-separable P is in particular both multilinear and homogeneous of degree d . Degree-separability enables a trade-off between the size of K and the query time for evaluation:

► **Theorem 2.** *Let s divide d and d/s divide $q-1$. There is a set $K \subseteq \mathbb{F}_q^n$ of size $|K| \leq (d/s+1)^{n+s}$ along with $g_1, g_2, \dots, g_{(q-1)^s} : \mathbb{F}_q^n \rightarrow K$ and $\gamma_1, \gamma_2, \dots, \gamma_{(q-1)^s} \in \mathbb{F}_q$ such that for every degree-separable degree- d $P \in \mathbb{F}_q[x]$ relative to a fixed partition X_1, X_2, \dots, X_d and every vector $a \in \mathbb{F}_q^n$,*

$$P(a) = \sum_{j=1}^{(q-1)^s} \gamma_j P(g_j(a)).$$

¹ In contrast, Kedlaya and Umans [11] focus on the case $n \leq d^{o(1)}$; cf. [11, Corollaries 4.3, 4.5, and 6.4]. *Notational caveat:* Kedlaya and Umans use “ m ” for the number of variables.

² For example, $M(q) = O((\log q)^{1+\epsilon})$ holds for any constant $\epsilon > 0$; we refer to e.g. von zur Gathen and Gerhard [24] for sharper bounds.

► **Remark.** The construction in Theorem 2 is explicit in the sense that (a) there is an algorithm that in time $O(|K|nqM(q))$ lists the elements of K ; and (b) there is an algorithm that in time $O(n(q-1)^s sqM(q))$ computes the values $g_j(a) \in \mathbb{F}_q^n$ and $\gamma_j \in \mathbb{F}_q$ for all $j = 1, 2, \dots, (q-1)^s$ when given $a \in \mathbb{F}_q^n$ as input.

We need K to contain enough points that “interpolation” at all the other points is possible. One intuition for designing a small $K \subseteq \mathbb{F}_q^n$ for polynomial evaluation is that such a set must enable “localization” of any target polynomial inside the set. At one extreme, we may think of the simplest non-constant family of polynomials, namely *lines*. In Euclidean spaces, this line of thought leads to the study of dimensionality of sets that contain a unit line segment in every direction, or the *Keakeya problem*, which has been extensively studied since the 1920s and the seminal work of Besicovitch [2]. We refer to Wolff [26], Mockenhaupt and Tao [18], and Dvir [9, 10] for surveys both in the continuous and finite settings. In what follows we focus on finite vector spaces.

► **Definition 3.** A *Keakeya set* (or *Besicovitch set*) in a vector space of dimension n over \mathbb{F}_q is a subset $K \subseteq \mathbb{F}_q^n$ together with a function $f : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^n$ such that for every vector $a \in \mathbb{F}_q^n$ and every scalar $\tau \in \mathbb{F}_q$ it holds that

$$f(a) + \tau a \in K. \tag{1}$$

That is, a Keakeya set K has the property that for any possible direction of a line in \mathbb{F}_q^n (that is, any nonzero vector $a \in \mathbb{F}_q^n$), the set K contains an entire line (through $f(a)$) with this direction. To support our objective of polynomial evaluations for higher-order curves than lines, an intuition is now to generalize (1) to polynomials of higher degree in the indeterminate τ . This is the methodological gist of our main contribution in this paper, which will be described further in Section 2.

As an illustrative application of our new data structures, we use Theorem 2 to derive a faster algorithm for computing fermionants, which are a family of self-reducible and degree-separable polynomials introduced by Chandrasekharan and Wiese [7] to generalize various fundamental polynomials. We start with a brief introduction to fermionants to motivate their study from a computational perspective.

1.2 Fermionants

We continue to work over \mathbb{F}_q . As usual, S_m is the symmetric group over $[m] = \{1, 2, \dots, m\}$. We write $c(\sigma)$ for the number of cycles in a permutation $\sigma \in S_m$, where each fixed point of σ is counted as a cycle of length 1. Let $A = (a_{ij} : i, j \in [m])$ be an $m \times m$ matrix of indeterminates. The *fermionant* of A with (indeterminate) parameter t is the $(m^2 + 1)$ -variable polynomial

$$\text{fer}_t A = (-1)^m \sum_{\sigma \in S_m} (-t)^{c(\sigma)} \prod_{i=1}^m a_{i, \sigma(i)}. \tag{2}$$

The fermionant is multilinear and homogeneous of degree m with respect to the variables $\{a_{i,j}\}$, and of degree m with respect to t . Furthermore, note that with respect to $\{a_{i,j}\}$ the fermionant is degree-separable under the partition $\{\{a_{i,j} : j \in [m]\} : i \in [m]\}$.

The fermionant captures several extensively studied algebraic and combinatorial invariants, such as the determinant of a matrix

$$\det A = (-1)^m \sum_{\sigma \in S_m} (-1)^{c(\sigma)} \prod_{i=1}^m a_{i, \sigma(i)},$$

the permanent of a matrix

$$\text{per } A = \sum_{\sigma \in S_m} \prod_{i=1}^m a_{i, \sigma(i)},$$

the generating function for directed Hamiltonian cycles

$$\text{hc } A = \sum_{\substack{\sigma \in S_m \\ c(\sigma)=1}} \prod_{i=1}^m a_{i, \sigma(i)},$$

as well as certain partition functions of strongly correlated electron systems in statistical physics (see Chandrasekharan and Wiese [7]). It is immediate that the aforementioned invariants can be obtained as special cases of the fermionant via

$$\det A = \text{fer}_1 A, \quad \text{per } A = (-1)^m \text{fer}_{-1} A, \quad \text{and} \quad \text{hc } A = (-1)^{m-1} \{t\} \text{fer}_t A,$$

where in the last equality we write $\{t^k\}P$ for the coefficient of t^k in the polynomial P .

The invariants captured by the fermionant have received such substantial attention that it is not possible to discuss the literature exhaustively here. For example, the permanent and the determinant are central to arithmetic circuit complexity [22] and geometric complexity theory [15]. Similarly, the numerous symmetries and self-reducibility properties of fermionants enable their use in e.g. interactive proof systems [5, 16, 25]. We restrict our present discussion of earlier work mostly to algorithms for the permanent.

Computing the permanent of a given $m \times m$ matrix appears to be an extremely hard problem. Indeed, the best known general algorithm is over 50 years old, given by Ryser [19] in 1963, and it uses $O(2^m m)$ arithmetic operations. Valiant [23] proved that the permanent for $\{0, 1\}$ -matrix inputs is $\#P$ -hard, even if the number of ones per row is at most three. In the more general setting of fermionants, Mertens and Moore [17] showed that the fermionant is $\#P$ -hard for any $\tau > 2$ and $\oplus P$ -hard for $\tau = 2$, even for the adjacency matrices of planar graphs. For the permanent, no less-than- 2^m -sized arithmetic circuit is known despite substantial efforts (for example, it is a prominent open problem in the *Art of Computer Programming* [12]).

However, there are faster ways to compute the permanent if we allow random-access tabulation *along with* arithmetic operations. Most notably, there are modest speed-ups for $\{0, 1\}$ -matrices over the integers. Bax and Franklin [1] gave an $2^{m-\Omega(m^{1/3}/\log m)}$ expected time algorithm. Recently, Björklund [3] presented a deterministic $2^{m-\Omega(\sqrt{m/\log q})}$ time algorithm over any finite field of order $q \geq m^2 + 1$, by exploiting the self-reducibility of the permanent. Applying the Chinese Remainder Theorem, he also obtains a $2^{m-\Omega(\sqrt{m/\log m})}$ -time algorithm for integer matrices with entries whose absolute value is bounded from above by a constant. There are also faster algorithms for sparse matrices. Cygan and Pilipczuk [8] gave a $2^{m-\Omega(m/r)}$ time algorithm for matrices with at most r non-zero entries per row. Very recently, Björklund, Husfeldt, and Lyckberg [4] and Björklund, Kaski, and Koutis [6] show that if the result is bounded in absolute value by c^m for a constant $c > 1$, then there are $2^{m(1-1/c^{\Omega(1)})} m^{O(1)}$ -time algorithms for the permanent and the number of directed Hamiltonian cycles, respectively. Both algorithms work by computing the permanent and the number of directed Hamiltonian cycles modulo small primes. In particular, the algorithms over \mathbb{F}_p run in time $2^{m(1-1/p^{\Omega(1)})}$, faster than the algorithms of this paper for small p .

Our main technical result for fermionants is that, given mild technical conditions on the order of the field, we can compute obtain a faster algorithm over finite fields:

► **Theorem 4.** *There is an algorithm that computes the fermionant $\text{fer}_t A \in \mathbb{F}_q[t]$ of a given $m \times m$ matrix A with entries in \mathbb{F}_q in time $2^{m-\Omega(\sqrt{m/\log \log q})} O(M(q))$, provided that $q - 1$ has a divisor in the interval $(1.1 \log q, 10 \log q)$, $q \geq m^2 + 1$, and $m = \omega(\log^2 q \log \log q)$.*

The Chinese Remainder Theorem and a uniform variant of the Prime Number Theorem for arithmetic progressions yield the following corollary for integer-valued fermionants.

► **Corollary 5.** *Let τ be an integer with $|\tau| \leq O(m)$ and let M be a constant. The fermionant $\text{fer}_\tau A$ can be computed in time $2^{m-\Omega(\sqrt{m/\log \log m})}$, for all $m \times m$ matrices A with integer values in $[-M, M]$.*

The idea behind Theorem 4 is to apply our polynomial evaluation results to a self-reduction for fermionants. Following Björklund’s results for the permanent [3], we show how to compute a fermionant on an $m \times m$ matrix via $2^{m-k} m^{O(1)}$ calls to the fermionant on $k \times k$ matrices. Applying Theorem 2, we set k so that it is possible to evaluate the $k \times k$ fermionant polynomial over all points of K in $2^{0.999m}$ time. Once we know the polynomial on all points in K , we can then evaluate the fermionant on any $m \times m$ matrix in time about $2^{m-\Omega(k)} m^{O(1)}$. We show $k \approx \sqrt{m/\log \log q}$ suffices.

Organisation. In Section 2, we present our generalization of Kakeya sets in finite vector spaces, together with explicit constructions. Next in Section 3 we prove our main evaluation theorems, Theorem 1 and Theorem 2. In Section 4 we use the self-reducibility of the fermionant to prove Theorem 4 and Corollary 5, showing how to compute fermionants faster.

2 Generalized Kakeya sets in finite vector spaces

Here we study the following generalization of Kakeya sets for lines (Definition 3) to higher-degree polynomial curves:

► **Definition 6.** A *Kakeya set of degree r* in a vector space of dimension n over \mathbb{F}_q consists of a set $K \subseteq \mathbb{F}_q^n$ together with functions $f_0, f_1, \dots, f_{r-1} : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^n$ such that for every vector $a \in \mathbb{F}_q^n$ and every scalar $\tau \in \mathbb{F}_q$ it holds that

$$F(a, \tau) = f_0(a) + f_1(a)\tau + f_2(a)\tau^2 + \dots + f_{r-1}(a)\tau^{r-1} + a\tau^r \in K. \tag{3}$$

We say that a construction for Kakeya sets is *explicit* if

- (i) there is an algorithm that outputs K (given q, r , and n) in $O(|K|nrM(q))$ time, and
- (ii) there is an algorithm that given $a \in \mathbb{F}_q^n$ outputs the values $f_0(a), f_1(a), \dots, f_{r-1}(a) \in \mathbb{F}_q^n$ in $O(nrM(q))$ time.

The following construction of sparse Kakeya sets of degree r generalizes the design of the best known Kakeya sets (cf. Mockenhaupt and Tao [18], Saraf and Sudan [20], Dvir [9, §2.4], Kopparty, Lev, Saraf, and Sudan [13], and Kyureghyan, Müller, and Wang [14]).

► **Lemma 7.** *For every $r + 1$ that divides $q - 1$ there is an explicit Kakeya set $K \subseteq \mathbb{F}_q^n$ of degree r and size $|K| \leq \left(\frac{q-1}{r+1} + 1\right)^{n+1}$.*

Proof. We begin with three simple observations. First, since $r + 1$ divides $q - 1$, we have that $r + 1$ has a multiplicative inverse in \mathbb{F}_q .³ Second, for all $\alpha, \tau \in \mathbb{F}_q$ from the Binomial

³ Indeed, $q = p^a$ for a prime p and positive integer a . Note $r + 1$ has a multiplicative inverse if and only if p does not divide $r + 1$. By assumption we have $(r + 1)Q = p^a - 1$ for an integer Q and thus $r + 1 = pb$ for an integer b would lead to a contradiction $p(bQ - p^{a-1}) = -1$.

Theorem we have

$$\left(\frac{\alpha}{r+1} + \tau\right)^{r+1} - \tau^{r+1} = \sum_{i=0}^{r-1} \binom{r+1}{i} \left(\frac{\alpha}{r+1}\right)^{r+1-i} \tau^i + \alpha\tau^r. \quad (4)$$

Third, since the multiplicative subgroup \mathbb{F}_q^\times is cyclic of order $q-1$, the subgroup consisting of $(r+1)$ th powers of elements of \mathbb{F}_q^\times has size exactly $\frac{q-1}{r+1}$. Including the zero element, we observe that $|\{\beta^{r+1} : \beta \in \mathbb{F}_q\}| = \frac{q-1}{r+1} + 1$.

Let us now define $K \subseteq \mathbb{F}_q^n$ to consist of all vectors of the form

$$\left(\left(\frac{\alpha_1}{r+1} + \tau\right)^{r+1} - \tau^{r+1}, \left(\frac{\alpha_2}{r+1} + \tau\right)^{r+1} - \tau^{r+1}, \dots, \left(\frac{\alpha_n}{r+1} + \tau\right)^{r+1} - \tau^{r+1}\right) \quad (5)$$

with $\alpha_1, \alpha_2, \dots, \alpha_n, \tau \in \mathbb{F}_q$. It follows immediately from (5) and our third observation that $|K| \leq \left(\frac{q-1}{r+1} + 1\right)^{n+1}$. Furthermore, (4) and (5) imply that the generalized Kakeya property (3) holds when we define the functions $f_i : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^n$ for all $i = 0, 1, \dots, r-1$ and $a = (\alpha_1, \alpha_2, \dots, \alpha_n) \in \mathbb{F}_q^n$ by

$$f_i(a) = \left(\binom{r+1}{i} \left(\frac{\alpha_1}{r+1}\right)^{r+1-i}, \binom{r+1}{i} \left(\frac{\alpha_2}{r+1}\right)^{r+1-i}, \dots, \binom{r+1}{i} \left(\frac{\alpha_n}{r+1}\right)^{r+1-i}\right). \quad (6)$$

It is immediate from the definitions (5) and (6) that the construction is explicit. ◀

3 Polynomial evaluation

This section proves our two main theorems for polynomial evaluation. The key idea is Mellin-transform-like sieving (8) enabled by an elementary observation about sums over finite fields (7) below, which we then extend to an s -fold product form in (12).

Let us start with a homogeneous version of Theorem 1.

► **Lemma 8.** *Let d divide $q-1$. There is a set $K \subseteq \mathbb{F}_q^n$ of size $|K| \leq (d+1)^{n+1}$ together with functions $g_1, g_2, \dots, g_{q-1} : \mathbb{F}_q^n \rightarrow K$ and coefficients $\gamma_1, \gamma_2, \dots, \gamma_{q-1} \in \mathbb{F}_q$ such that for every homogeneous polynomial $P \in \mathbb{F}_q[x]$ of degree $h \leq d$ and every vector $a \in \mathbb{F}_q^n$,*

$$P(a) = \sum_{j=1}^{q-1} \gamma_j P(g_j(a)).$$

Proof. Let d divide $q-1$. Set $r = (q-1)/d - 1$, and note that $r+1$ divides $q-1$. Apply Lemma 7 to obtain K and the functions f_0, f_1, \dots, f_{r-1} . Let $P \in \mathbb{F}_q[x]$ be a homogeneous polynomial of degree $h \leq d$ over the indeterminates $x = (x_1, x_2, \dots, x_n)$, and let $a = (\alpha_1, \alpha_2, \dots, \alpha_n) \in \mathbb{F}_q^n$ be an assignment of values to the indeterminates. Our goal is to compute the value $P(a) \in \mathbb{F}_q$ using evaluations of P at K . Recalling the function $F(a, \tau)$ from (3), we will rely on values of the composition $P(F(a, \tau))$ for $\tau \in \mathbb{F}_q$ to obtain $P(a)$.

Towards this end, we first observe that

$$\sum_{\tau \in \mathbb{F}_q^\times} \tau^e = \begin{cases} -1 & \text{if } q-1 \text{ divides } e, \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

To see this, let g be a generator of the multiplicative subgroup \mathbb{F}_q^\times . If $q-1$ divides e then $\tau^e = 1$ for all τ , and thus the sum is $|\mathbb{F}_q^\times| = q-1$ (modulo the characteristic). Otherwise, $g^e \neq 1$, and we have $\sum_{\tau \in \mathbb{F}_q^\times} \tau^e = \sum_{\tau \in \mathbb{F}_q^\times} (g\tau)^e = g^e \sum_{\tau \in \mathbb{F}_q^\times} \tau^e$, so the sum must be 0.

Let $t = q - 1 - rh$ and observe that $t \geq 1$. We now claim that

$$P(a) = - \sum_{\tau \in \mathbb{F}_q^\times} \tau^t P(F(a, \tau)). \tag{8}$$

By linearity, it suffices to consider the case when P is a single monomial $P = x_1^{h_1} x_2^{h_2} \dots x_n^{h_n}$ of degree $h = h_1 + h_2 + \dots + h_n \leq d$. Recalling (3) and (7), we observe that the right-hand side of (8) expands to

$$\begin{aligned} & - \sum_{\tau \in \mathbb{F}_q^\times} \tau^t P(F(a, \tau)) \\ &= - \sum_{\tau \in \mathbb{F}_q^\times} \tau^{q-1-rh} \left(\tau^{rh} \alpha_1^{h_1} \alpha_2^{h_2} \dots \alpha_n^{h_n} + \tau^{rh-1}(\dots) + \tau^{rh-2}(\dots) + \dots + \tau^0(\dots) \right) \\ &= - \sum_{\tau \in \mathbb{F}_q^\times} \left(\tau^{q-1} \alpha_1^{h_1} \alpha_2^{h_2} \dots \alpha_n^{h_n} + \tau^{q-2}(\dots) + \tau^{q-3}(\dots) + \dots + \tau^{q-1-rh}(\dots) \right) \\ &= \alpha_1^{h_1} \alpha_2^{h_2} \dots \alpha_n^{h_n} \\ &= P(a). \end{aligned}$$

That is, by multiplying each term by τ^t , we ensure that all other terms appearing inside of $P(F(a, \tau))$ cancel, except for the desired term $\alpha_1^{h_1} \alpha_2^{h_2} \dots \alpha_n^{h_n}$ which is the coefficient of τ^{rh} .

Now let $\beta_1, \beta_2, \dots, \beta_{q-1}$ be an enumeration of the elements of \mathbb{F}_q^\times . For all $j = 1, 2, \dots, q-1$, set $g_j(a) = F(a, \beta_j)$ and $\gamma_j = -\beta_j^t$. The lemma now follows from (8). ◀

3.1 Proof of Theorem 1

We are now ready to prove Theorem 1. Our strategy is to interpolate the homogeneous components of our given polynomial, then apply Lemma 8. Towards this end, let $P \in \mathbb{F}_q[x]$ have degree at most d and let $P = \sum_{h=0}^d P_h$ where $P_h \in \mathbb{F}_q[x]$ is either zero or homogeneous of degree h , for all $h = 0, 1, \dots, d$. Let $\nu_0, \nu_1, \dots, \nu_d$ be any $d+1$ distinct elements of \mathbb{F}_q . Recalling the definition of K in (5), let $\hat{K} \subseteq \mathbb{F}_q^n$ be the set of all vectors of the form

$$\nu \left(\left(\frac{\alpha_1}{r+1} + \tau \right)^{r+1} - \tau^{r+1}, \left(\frac{\alpha_2}{r+1} + \tau \right)^{r+1} - \tau^{r+1}, \dots, \left(\frac{\alpha_n}{r+1} + \tau \right)^{r+1} - \tau^{r+1} \right) \tag{9}$$

where $\alpha_1, \alpha_2, \dots, \alpha_n, \tau \in \mathbb{F}_q$, and $\nu \in \{\nu_0, \nu_1, \dots, \nu_d\}$.

In particular, from (9) and (5) we have that $|\hat{K}| \leq (d+1)|K|$.

Assuming we have constant-time access to $P(a)$ for all $a \in \hat{K}$, we can access each P_h at $k \in K$ by univariate interpolation over the $d+1$ distinct values of ν , via the identity $P(\nu k) = \sum_{h=0}^d P_h(k) \nu^h$. That is, for $h, j = 0, 1, \dots, d$, let $\lambda_{hj} \in \mathbb{F}_q$ be the Lagrange interpolation coefficients that satisfy $P_h(k) = \sum_{j=0}^d \lambda_{hj} P(\nu_j k)$ for all $k \in K$. Observe in particular that the coefficients λ_{hj} depend only on $\nu_0, \nu_1, \dots, \nu_d$. With access to values of P_h at K , given a query $a \in \mathbb{F}_q^n$ we can use Lemma 8 to sieve for $P_h(a)$ for each $h = 0, 1, \dots, d$. That is, we have $P(a) = \sum_{h=0}^d P_h(a) = - \sum_{h=0}^d \sum_{\tau \in \mathbb{F}_q^\times} \sum_{j=0}^d \tau^{q-1-rh} \lambda_{hj} P(\nu_j F(a, \tau))$. ◀

3.2 Proof of Theorem 2

Suppose s divides d and d/s divides $q-1$. Let X_1, X_2, \dots, X_d be the partition of variables for degree-separability. For $i = 1, 2, \dots, s$, take

$$Y_i = X_{(i-1)d/s+1} \cup X_{(i-1)d/s+2} \cup \dots \cup X_{id/s}$$

and observe that $|Y_i| = n/s$ for all i . Furthermore, observe that every monomial of a polynomial $P \in \mathbb{F}_q[x]$ that is degree-separable relative to X_1, X_2, \dots, X_d has degree exactly d/s when restricted to the variables of Y_i .

Let us extend the construction in Lemma 7 into an s -fold product form over the partition Y_1, Y_2, \dots, Y_s . Accordingly, we work with a multivariate polynomial over s indeterminates $\tau_1, \tau_2, \dots, \tau_s$ instead of a univariate polynomial (3) over τ . Let $a = (\alpha_1, \alpha_2, \dots, \alpha_n) \in \mathbb{F}_q^n$ and let us write $a_{Y_i} \in \mathbb{F}_q^{n_s/d}$ for the restriction of a to coordinates in Y_i . Set $r = (q - 1)s/d - 1$. Let us write $F_{Y_i}(a_{Y_i}, \tau_i) \in \mathbb{F}_q^n$ for the vector obtained by applying the construction given by (3) and (6) to the vector a_{Y_i} and τ_i , thereby obtaining a vector of length ns/d indexed by Y_i , followed by padding with 0-entries outside the indices Y_i to obtain a vector of length n . Let us now define the (vector-valued) multivariate polynomial

$$F(a, \tau_1, \tau_2, \dots, \tau_s) = F_{Y_1}(a_{Y_1}, \tau_1) + F_{Y_2}(a_{Y_2}, \tau_2) + \dots + F_{Y_s}(a_{Y_s}, \tau_s). \tag{10}$$

We observe by (3), (6), and (4) that $F(a, \tau_1, \tau_2, \dots, \tau_s)$ ranges over all vectors of the form

$$\begin{aligned} & \left(\left(\frac{\alpha_1}{r+1} + \tau_1 \right)^{r+1} - \tau_1^{r+1}, \left(\frac{\alpha_2}{r+1} + \tau_1 \right)^{r+1} - \tau_1^{r+1}, \dots, \left(\frac{\alpha_{n/s}}{r+1} + \tau_1 \right)^{r+1} - \tau_1^{r+1}, \right. \\ & \left(\frac{\alpha_{n/s+1}}{r+1} + \tau_2 \right)^{r+1} - \tau_2^{r+1}, \left(\frac{\alpha_{n/s+2}}{r+1} + \tau_2 \right)^{r+1} - \tau_2^{r+1}, \dots, \left(\frac{\alpha_{2n/s}}{r+1} + \tau_2 \right)^{r+1} - \tau_2^{r+1}, \\ & \dots, \\ & \left. \left(\frac{\alpha_{n-n/s+1}}{r+1} + \tau_s \right)^{r+1} - \tau_s^{r+1}, \left(\frac{\alpha_{n-n/s+2}}{r+1} + \tau_s \right)^{r+1} - \tau_s^{r+1}, \dots, \left(\frac{\alpha_n}{r+1} + \tau_s \right)^{r+1} - \tau_s^{r+1} \right) \end{aligned} \tag{11}$$

with $\alpha_1, \alpha_2, \dots, \alpha_n, \tau_1, \tau_2, \dots, \tau_s \in \mathbb{F}_q$. We define K to be the set of all such vectors. By similar reasoning as in the proof of Theorem 1, note that $|K| \leq \left(\frac{q-1}{r+1} + 1 \right)^{n+s} = \left(\frac{d}{s} + 1 \right)^{n+s}$.

Let $t = q - 1 - rd/s$ and observe that $t \geq 1$. From (7) and proceeding analogously as with the reasoning for (8) in the proof of Theorem 1, we thus have

$$P(a) = (-1)^s \sum_{\tau_1, \tau_2, \dots, \tau_s \in \mathbb{F}_q^\times} \tau_1^t \tau_2^t \dots \tau_s^t P(F(a, \tau_1, \tau_2, \dots, \tau_s)). \tag{12}$$

Let $\beta_1, \beta_2, \dots, \beta_{q-1}$ be an enumeration of the elements of \mathbb{F}_q^\times . For all $j = (j_1, j_2, \dots, j_s) \in \{1, 2, \dots, q-1\}^s$ take

$$g_j(a) = F(a, \beta_{j_1}, \beta_{j_2}, \dots, \beta_{j_s}) \quad \text{and} \quad \gamma_j = (-1)^s \beta_{j_1}^t \beta_{j_2}^t \dots \beta_{j_s}^t.$$

The theorem now follows from (12). ◀

4 Fermionants

This section proves our two main theorems for evaluating fermionants. We start by noting that the fermionant is self-reducible, a result that easily follows from earlier work by Björklund [3], followed by the proofs of our present main theorems.

4.1 Self-reducibility of the fermionant

This subsection reviews how Björklund’s [3] self-reducibility for permanents can be extended to fermionants. In essence, his methodology can be used to reduce the task of computing one fermionant of size $m \times m$ to the task of computing $2^{m-k} m^{O(1)}$ fermionants of size $k \times k$. We

stress that this subsection is provided for ease of exposition only and no claim of originality is made.

Let r, t , and a_{ij} for $i, j \in [m]$ be polynomial indeterminates and let \mathbb{F} be the coefficient field. For $S \subseteq [m]$, $i, j \in S$, and $\ell = 0, 1, \dots, m$, consider the inductively-defined family of polynomials:

$$W_{\ell,i,j}^S(r) = \begin{cases} 1 & \text{if } \ell = 0 \text{ and } i = j; \\ 0 & \text{if } \ell = 0 \text{ and } i \neq j; \\ \sum_{u \in S} a_{iu} r W_{\ell-1,u,j}^S(r) & \text{if } \ell \geq 1. \end{cases} \tag{13}$$

The polynomial $W_{\ell,i,j}^S$ can be viewed as a multivariate generating function for (edge-multisets of) walks of length ℓ inside $S \subseteq [m]$ that start at i and end at j , on the complete graph of m vertices. The degree of each monomial in the indeterminate r is equal to ℓ . The indeterminates a_{uv} track the edges (u, v) traversed by the walk, with degree indicating the multiplicity, that is, how many times each edge was traversed.

For $S \subseteq [m]$ and $i \in [m]$, let us write $S_{\geq i} = \{u \in S : i \leq u\}$. For $S \subseteq [m]$, $i \in S$, and $\ell = 0, 1, \dots, m$, introduce the following bivariate polynomial

$$C_i^S(r, t) = 1 - t \sum_{\ell=1}^m W_{\ell,i,i}^{S_{\geq i}}(r). \tag{14}$$

This polynomial forms a multivariate generating function for (edge-multisets of) closed walks inside S and “anchored” at i , including the possibility of no walk at all. Here, by “anchored” at i we mean that the lowest-numbered vertex of the closed walk is i . Let

$$C^S(r, t) = \prod_{i \in S} C_i^S(r, t). \tag{15}$$

Let $k = 0, 1, \dots, m$. For $i, j \in [k]$ and $S \subseteq [m] \setminus [k]$, introduce the univariate polynomial

$$\tilde{a}_{i,j}^S(r) = a_{ij} + \sum_{\ell=0}^{m-1} \sum_{u,v \in S} a_{iu} W_{\ell,u,v}^S a_{vj} r. \tag{16}$$

This polynomial is a multivariate generating function for representing the (edge-multisets of) walk segments that traverse $[m]$ so that the first vertex of the segment is i and the last vertex of the segment is j ; the walk can either proceed directly from i to j , or perform a walk of length ℓ in S before ending at j . Let us arrange the coefficients $\tilde{a}_{i,j}^S(r)$ into a $k \times k$ matrix $\tilde{A}^S(r)$.

For a polynomial P in the indeterminate r , let us write $\{r^j\}P$ for the coefficient (polynomial) of the monomial r^j . By the principle of inclusion and exclusion, we have:

► **Theorem 9.** *For all $k = 0, 1, \dots, m$, we have the polynomial identity*

$$\text{fer}_t A = \{r^{m-k}\} \sum_{S \subseteq [m] \setminus [k]} (-1)^{|S|} C^S(r, t) \text{fer}_t \tilde{A}^S(r). \tag{17}$$

Observing that the right-hand side of (17) has degree at most m^2 in r , and using Lagrange interpolation together with dynamic programming on the recurrences (13), (14), (15), and (16), we have:

► **Theorem 10.** *Suppose $|\mathbb{F}| \geq m^2 + 1$ and let $k = 0, 1, \dots, m$. Then, there is a value $L = 2^{m-k} m^{O(1)}$ computable in time polynomial in m , and an algorithm that given as input*

7:10 Generalized Kakeya sets and faster fermionants

a matrix $A \in \mathbb{F}^{m \times m}$, $\tau \in \mathbb{F}$, and an integer $j = 1, 2, \dots, L$, runs in time $m^{O(1)}$, executes $m^{O(1)}$ arithmetic operations in \mathbb{F} , and outputs a matrix $\tilde{A}_j \in \mathbb{F}^{k \times k}$ together with a coefficient $\alpha_j \in \mathbb{F}$ such that:

$$\text{fer}_\tau A = \sum_{j=1}^L \alpha_j \text{fer}_\tau \tilde{A}_j. \quad (18)$$

In particular, the fermionant $\text{fer}_\tau A$ of a given $A \in \mathbb{F}^{m \times m}$ at $\tau \in \mathbb{F}$ can be computed in $2^m m^{O(1)}$ time and arithmetic operations in \mathbb{F} .

4.2 Proof of Theorem 4

Let $A \in \mathbb{F}_q^{m \times m}$ be given together with $\tau \in \mathbb{F}_q$. We seek to compute $\text{fer}_\tau A$ and will deploy the self-reducibility enabled by Theorem 10 towards this end. By assumption we have that $q-1$ has a divisor u with $1.1 \log q \leq u \leq 10 \log q$. Since $m = \omega(\log^2 q \log \log q)$, for all large enough m we can let k be a multiple of u with

$$0.98 \sqrt{m / \log \log q} \leq k \leq 0.99 \sqrt{m / \log \log q}.$$

With the objective of applying Theorem 2, take $n = k^2$, $d = k$, and $s = k/u$. Observe that the fermionant (2) of a $k \times k$ matrix A at $\tau \in \mathbb{F}_q$ is a degree-separable polynomial P of degree d over the n variables in A . Furthermore, s divides d and d/s divides $q-1$, so the assumptions of Theorem 2 hold. By Theorem 10 we can evaluate this P at any given point (that is, for any given $k \times k$ matrix) in time $2^k k^{O(1)}$ and operations in \mathbb{F}_q . The tabulation of P for Theorem 2 thus can be done in time

$$\begin{aligned} 2^k k^{O(1)} \left(\frac{d}{s} + 1 \right)^{n+s} M(q) &\leq 2^k k^{O(1)} (u+1)^{0.99m / \log \log q + \sqrt{m}} M(q) \\ &\leq 2^k k^{O(1)} (20 \log q)^{0.999m / \log \log q} M(q) \\ &\leq 2^{0.9999m} M(q). \end{aligned}$$

Once the tabulation of P is complete, we can use the algorithms in Theorem 2 to query the $2^{m-k} m^{O(1)}$ fermionants of size $k \times k$ required by (18) in time $O(n(q-1)^s s M(q))$ per query. Thus, the total time is at most

$$\begin{aligned} 2^{m-k} q^s m^{O(1)} M(q) &\leq 2^{m-0.98 \sqrt{m / \log \log q}} 2^{(\log q) 0.99 \sqrt{m / \log \log q} / (1.1 \log q)} m^{O(1)} M(q) \\ &\leq 2^{m-0.07 \sqrt{m / \log \log q}} m^{O(1)} M(q). \quad \blacktriangleleft \end{aligned}$$

4.3 Proof of Corollary 5

Here we show how to extend the algorithm to integers, via the Chinese Remainder Theorem. Let A be an integer matrix of size $m \times m$ with entries in $[-M, M]$ for $M = O(1)$. Let τ be an integer with $|\tau| = O(m)$. By Bertrand's postulate (e.g. [21, §I.1]) for all large enough m we can select a prime u with $5 \log m \leq u \leq 10 \log m$. Let us study the number of primes p in the interval $Mm^2 < p < Mm^4$ such that u divides $p-1$. Let us write φ for Euler's totient function and recall the uniform variant of the Prime Number Theorem for arithmetic progressions [21, Corollary 8.31]. Namely, there is a constant $\gamma > 0$ such that, for any function $h(x)$ tending to infinity with x , and uniformly for $x \geq 3$ and $1 \leq u \leq (\ln x)^2 / (h(x)^2 (\ln \ln x)^6)$, we have

$$\sum_{\substack{p \leq x \\ p \equiv 1 \pmod{u}}} 1 = \frac{x}{\varphi(u) \ln x} \left(1 + O\left(\frac{1}{(\ln x)^{\gamma h(x)}} \right) \right). \quad (19)$$

Here the left-hand side sum in (19) is over all primes p at most x congruent to 1 modulo u .

Since u is prime, we have $\varphi(u) = u - 1 = \Theta(\log m)$. Thus from (19) we conclude that for all large enough m there exist at least $2m$ distinct primes p such that both $Mm^2 < p < Mm^4$ and u divides $p - 1$. With the objective of satisfying the assumptions of Theorem 4, we conclude u is in the interval $(1.1 \log p, 10 \log p)$ for these $2m$ primes p . Indeed, since M is a constant, for all large enough m we have $1.99 \log m \leq \log p \leq 4.01 \log m$, which implies $(5/4.01) \log p \leq 5 \log m \leq u \leq 10 \log m \leq (10/1.99) \log p$.

From (2) we observe that $|\text{fer}_\tau A| \leq m! \cdot O(m)^m M^m < \frac{1}{2} m^{4m} M^{2m}$. Applying the Chinese Remainder Theorem together with Theorem 4 on A and τ over \mathbb{F}_p for each of the $2m$ primes p in turn, we recover $\text{fer}_\tau A$ over the integers, in time $2^{m-\Omega(\sqrt{m/\log \log m})}$. ◀

Acknowledgements. This research was funded by the Swedish Research Council grant VR 2016-03855 “Algebraic Graph Algorithms” (A.B.), the European Research Council under the European Union’s Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement 338077 “Theory and Practice of Advanced Search and Enumeration” (P.K.), and the U.S. National Science Foundation under grants CCF-1741638 and CCF-1741615 (R.W.).

References

- 1 Eric T. Bax and Joel Franklin. A finite-difference sieve to count paths and cycles by length. *Inf. Process. Lett.*, 60(4):171–176, 1996. doi:10.1016/S0020-0190(96)00159-7.
- 2 A. S. Besicovitch. On Kakeya’s problem and a similar one. *Math. Z.*, 27(1):312–320, 1928. doi:10.1007/BF01171101.
- 3 Andreas Björklund. Below all subsets for some permutational counting problems. In Rasmus Pagh, editor, *15th Scandinavian Symposium and Workshops on Algorithm Theory, SWAT 2016, June 22-24, 2016, Reykjavik, Iceland*, volume 53 of *LIPICs*, pages 17:1–17:11. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016. doi:10.4230/LIPICs.SWAT.2016.17.
- 4 Andreas Björklund, Thore Husfeldt, and Isak Lyckberg. Computing the permanent modulo a prime power. *Inf. Process. Lett.*, 125:20–25, 2017. doi:10.1016/j.ipl.2017.04.015.
- 5 Andreas Björklund and Petteri Kaski. How proofs are prepared at Camelot: Extended abstract. In George Giakkoupis, editor, *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing, PODC 2016, Chicago, IL, USA, July 25-28, 2016*, pages 391–400. ACM, 2016. doi:10.1145/2933057.2933101.
- 6 Andreas Björklund, Petteri Kaski, and Ioannis Koutis. Directed Hamiltonicity and out-branchings via generalized Laplacians. In Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl, editors, *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July 10-14, 2017, Warsaw, Poland*, volume 80 of *LIPICs*, pages 91:1–91:14. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017. doi:10.4230/LIPICs.ICALP.2017.91.
- 7 Shailesh Chandrasekharan and Uwe-Jens Wiese. Partition functions of strongly correlated electron systems as “fermionants”. 2011. arXiv:1108.2461v1. URL: <https://arxiv.org/abs/1108.2461>.
- 8 Marek Cygan and Marcin Pilipczuk. Faster exponential-time algorithms in graphs of bounded average degree. In *Automata, Languages, and Programming - 40th International Colloquium, ICALP 2013, Riga, Latvia, July 8-12, 2013, Proceedings, Part I*, pages 364–375, 2013. doi:10.1007/978-3-642-39206-1_31.
- 9 Zeev Dvir. From randomness extraction to rotating needles. *Electronic Colloquium on Computational Complexity (ECCC)*, 16:77, 2009. URL: <http://ecc.eccc.hpi-web.de/report/2009/077>.

- 10 Zeev Dvir. Incidence theorems and their applications. *CoRR*, abs/1208.5073, 2012. URL: <http://arxiv.org/abs/1208.5073>.
- 11 Kiran S. Kedlaya and Christopher Umans. Fast polynomial factorization and modular composition. *SIAM J. Comput.*, 40(6):1767–1802, 2011. doi:10.1137/08073408X.
- 12 Donald E. Knuth. *The Art of Computer Programming, Volume 2: Seminumerical Algorithms*. Addison-Wesley, 3rd edition, 1998.
- 13 Swastik Kopparty, Vsevolod F. Lev, Shubhangi Saraf, and Madhu Sudan. Kakeya-type sets in finite vector spaces. *J. Algebraic Combin.*, 34(3):337–355, 2011. doi:10.1007/s10801-011-0274-8.
- 14 Gohar Kyureghyan, Peter Müller, and Qi Wang. On the size of Kakeya sets in finite vector spaces. *Electron. J. Combin.*, 20(3):#P36, 2013. URL: <http://www.combinatorics.org/ojs/index.php/eljc/article/view/v20i3p36>.
- 15 J. M. Landsberg. An introduction to geometric complexity theory. *Eur. Math. Soc. Newsl.*, 99:10–18, 2016.
- 16 Carsten Lund, Lance Fortnow, Howard J. Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *J. ACM*, 39(4):859–868, 1992. doi:10.1145/146585.146605.
- 17 Stephan Mertens and Cristopher Moore. The complexity of the fermionant, and immanants of constant width. 2011. arXiv:1110.1821. URL: <https://arxiv.org/abs/1110.1821>.
- 18 Gerd Mockenhaupt and Terence Tao. Restriction and Kakeya phenomena for finite fields. *Duke Math. J.*, 121(1):35–74, 2004. doi:10.1215/S0012-7094-04-12112-8.
- 19 H. J. Ryser. *Combinatorial Mathematics*. Number 14 in The Carus Mathematical Monographs. Mathematical Association of America, 1963.
- 20 Shubhangi Saraf and Madhu Sudan. An improved lower bound on the size of Kakeya sets over finite fields. *Anal. PDE*, 1(3):375–379, 2008. doi:10.2140/apde.2008.1.375.
- 21 Gérald Tenenbaum. *Introduction to Analytic and Probabilistic Number Theory*, volume 163 of *Graduate Studies in Mathematics*. American Mathematical Society, 3rd edition, 2015.
- 22 Leslie G. Valiant. Completeness classes in algebra. In Michael J. Fischer, Richard A. DeMillo, Nancy A. Lynch, Walter A. Burkhard, and Alfred V. Aho, editors, *Proceedings of the 11th Annual ACM Symposium on Theory of Computing, April 30 - May 2, 1979, Atlanta, Georgia, USA*, pages 249–261. ACM, 1979. doi:10.1145/800135.804419.
- 23 Leslie G. Valiant. The complexity of computing the permanent. *Theor. Comput. Sci.*, 8:189–201, 1979. doi:10.1016/0304-3975(79)90044-6.
- 24 Joachim von zur Gathen and Jürgen Gerhard. *Modern Computer Algebra*. Cambridge University Press, Cambridge, third edition, 2013. doi:10.1017/CB09781139856065.
- 25 Richard Ryan Williams. Strong ETH breaks with Merlin and Arthur: Short non-interactive proofs of batch evaluation. In Ran Raz, editor, *31st Conference on Computational Complexity, CCC 2016, May 29 to June 1, 2016, Tokyo, Japan*, volume 50 of *LIPICs*, pages 2:1–2:17. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016. doi:10.4230/LIPICs.CCC.2016.2.
- 26 Thomas Wolff. Recent work connected with the Kakeya problem. In *Prospects in Mathematics (Princeton, NJ, 1996)*, pages 129–162. Amer. Math. Soc., Providence, RI, 1999.