Algebraic Cayley Graphs over Finite Fields☆

M. Lu[a], D. Wan[b], L.-P. Wang[c,*], X.-D. Zhang[d]

[a] *Department of Mathematical Sciences, Tsinghua University, Beijing, China*
[b] *Department of Mathematics, University of California, Irvine, CA 92697-3875, USA*
[c] *Insitutute of Information Engineering, Chinese Academy of Sciences, Beijing, China*
[d] *Department of Mathematics, and MOE-LSC, Shanghai Jiao Tong University, Shanghai, China*

Abstract

A new algebraic Cayley graph is constructed using finite fields. It provides a more flexible source of expander graphs. Its connectedness, the number of connected components, and diameter bound are studied via Weil's estimate for character sums. Furthermore, we study the algorithmic problem of computing the number of connected components and establish a link to the integer factorization problem.

Keywords: Algebraic Cayley graphs, Character sums, Expander graphs

1. Introduction

For a subset S of a finite abelian group Γ, the Cayley graph $Cay(\Gamma, S)$ is the directed graph with vertex set Γ, and edge set $\{b_1 \to b_2 | b_1 - b_2 \in S\}$. Cayley graphs play a central role in the construction of expander graphs. A randomly chosen Cayley graph $Cay(\Gamma, S)$ often has good properties with nontrivial probability. However, deterministically constructing one such good graph is often more difficult. Typically one needs to assume additional structure on the group Γ and its subset S. By an algebraic Cayley graph, we mean that Γ is the multiplicative group of a finite commutative ring and $S \subset \Gamma$ is a subset with certain algebraic structure such as a box or an interval in some

☆Partially supported by 973 Program (2013CB834203) and National Natural Science Foundation of China (Grant No. 61373019, 61170289 and 11271256).

*Corresponding author

Email addresses: mlu@math.tsinghua.edu.cn (M. Lu), dwan@math.uci.edu (D. Wan), wangliping@iie.ac.cn (L.-P. Wang), xiaodong@sjtu.edu.cn (X.-D. Zhang)

sense. The box algebraic structure makes it possible to use powerful tools from number theory to prove conditionally (assuming some sort of Riemann hypothesis) that an algebraic Cayley graph $Cay(\Gamma, S)$ does have the desired properties if the box is suitably large. In this way, algebraic Cayley graphs provide a rich source of expander graphs.

An important example is given by Chung [1], who uses the multiplicative group of a finite extension of a finite field and takes the subset to be a line in certain sense. The advantage of working with a finite field is that the needed estimate can sometimes be proved using the celebrated Weil bound for curves over finite fields. In this paper, we introduce a more general construction using the multiplicative group of a finite field and taking the subset to be those elements represented by certain primary polynomials.

Let $\mathbb{F}_q$ be a finite field of q elements with characteristic p. Let $f(x)$ be an irreducible polynomial of degree $n > 1$ over $\mathbb{F}_q$. Our group Γ will be

$$\Gamma_f = (\mathbb{F}_q[x]/(f(x)))^* = (\mathbb{F}_q[\alpha])^* = \mathbb{F}_{q^n}^*, \quad \alpha = \overline{x}.$$

The group Γ_f is cyclic of order $q^n - 1$. A polynomial $g(x) \in \mathbb{F}_q[x]$ of degree $d > 0$ is called primary if $g(x)$ is a power of an irreducible polynomial. For $1 \leq d < n$, let P_d be the set of monic primary polynomials of degree d in $\mathbb{F}_q[x]$. Our subset S will be

$$E_d = \{g(\alpha) | g \in P_d\} \subset \Gamma_f.$$

Note that in the case $d = 1$, the subset $E_1 = \alpha + \mathbb{F}_q$ is a line in the n-dimensional $\mathbb{F}_q$-vector space $\mathbb{F}_{q^n}$.

Definition 1. *Let $G_d(n, q, \alpha)$ be the Cayley graph $Cay(\Gamma_f, E_d)$ with vertex set Γ_f and edge set $\{\beta_1 \to \beta_2 | \beta_2/\beta_1 \in E_d\}$.*

It is clear that $G_d(n, q, \alpha)$ is a regular directed graph of order $q^n - 1$ and its degree is given by

$$|E_d| = |P_d| = \sum_{k|d} \frac{1}{k} \sum_{s|k} \mu(s) q^{\frac{k}{s}} \sim \frac{q^d}{d},$$

where μ is the Möbius function. It should be noted that the graph $G_d(n, q, \alpha)$ depends not just on d, n, q but also on the choice of α (that is, the choice of the irreducible polynomial $f(x)$ which is used to present the extension field

$\mathbb{F}_{q^n}$). In the case $d = 1$, $G_1(n, q, \alpha)$ reduces to Chung's graph in [1], which has been studied extensively (see [2, 4, 5]). In this paper, we study the general d case. Our proof is more direct and uses Weil's bound for character sums.

Our first result is the following theorem.

Theorem 2. *Assume that $n < q^{d/2} + 1$. Then the graph $G_d(n, q, \alpha)$ is connected, and its diameter D satisfies the bound*

$$D \leq 2\frac{n}{d} + 1 + \frac{4\frac{n}{d}\log(n-1)}{d\log q - 2\log(n-1)}.$$

In the case $d = 1$, this reduces to the diameter bound in [1] and [5]. The above theorem gives a sufficient condition for the graph to be connected. If $n \geq q^{d/2} + 1$, the graph $G_d(n, q, \alpha)$ is not always connected, as the answer depends on the choice of α or the irreducible polynomial $f(x)$. More precisely, we have:

Theorem 3. *If $\ell > 1$ is a divisor of the integer $(q^n - 1)$ such that $n \geq 2d + 2(|P_d| + 1)\log_q \ell$, then there is at least one $\alpha \in \mathbb{F}_{q^n}$ of degree n such that the number of connected components of the graph $G_d(n, q, \alpha)$ is divisible by ℓ.*

If $q > 2$, $q^n - 1$ has the obvious divisor $(q-1) > 1$. We obtain the following result.

Corollary 4. *Assume that $q > 2$ and $n \geq 2d + 2(|P_d| + 1)$. Then there is at least one $\alpha \in \mathbb{F}_{q^n}$ of degree n such that the number of connected components of the graph $G_d(n, q, \alpha)$ is divisible by $(q-1)$. In particular, $G_d(n, q, \alpha)$ is not connected for at least one degree n element α.*

As $|P_d| \sim q^d/d$, the bound $2d + 2(|P_d| + 1) \sim 2q^d/d$ is roughly the square of the bound $q^{d/2}$ in Theorem 2. This shows that the condition in Theorem 2 is not too far from being sharp. For the remaining interval where

$$q^{d/2} + 1 \leq n \leq 2d + 2(|P_d| + 1) \sim 2q^d/d,$$

we have no results on the connectedness of the graph $G_d(n, q, \alpha)$. One does know the following crude combinatorial upper bound for the number $N_d(n, q, \alpha)$ of connected components of the graph $G_d(n, q, \alpha)$:

$$N_d(n, q, \alpha) \leq \frac{q^n - 1}{\binom{|P_d| + \lceil \frac{n}{d} \rceil}{\lceil \frac{n}{d} \rceil - 1}}.$$

For details, see Theorem 14 in section 2.

For a randomly chosen α, the graph $G_d(n, q, \alpha)$ is connected with non-negligible probability. To see this, we fix $g \in P_d$ to be monic irreducible. Then for a randomly chosen α, the element $g(\alpha)$ is a primitive root of $\mathbb{F}_{q^n}^*$ with probability

$$\frac{\phi(q^n - 1)}{q^n - 1} = \Omega(\frac{1}{\log\log(q^n - 1)}),$$

which is non-negligible, where ϕ is the Euler's totient function. Thus, the graph $G_d(n, q, \alpha)$ is connected with non-negligible probability. Unfortunately, constructing a primitive root (or even an element of high order) of any form is a well known difficult problem in computational number theory. In practical application, the difficulty is how to verify quickly that a given $G_d(n, q, \alpha)$ is connected and more generally, how to compute quickly the number of its connected components, using the sparse input size $(n \log q)^{O(1)}$ of the graph $G_d(n, q, \alpha)$. Ideally, we would like to have a deterministic algorithm with running time bounded by a polynomial in $(n \log q)^{O(1)}$, to compute the number of connected components. In this direction, we have the following conditional result.

Theorem 5. *Assume that the factorization of $q^n - 1$ is given. Then one can compute the number of connected components of $G_d(n, q, \alpha)$ in time $(n \log q)^{O(1)}$.*

It would be of great interest to remove the factorization assumption in the above theorem.

An important type of graphs is the so-called expander graph, which arises in questions about designing networks that connect many users while using a small number of switches. Expander graphs play an important role in computer science, mathematics, and the theory of communication networks. Please refer to the survey article [3]. A regular graph of degree k is an expander graph if the modulus of every nontrivial eigenvalue of the graph is much less than the trivial eigenvalue k. In the last section, we show that our graph $G_d(n, q, \alpha)$ provides a new source of expander graphs.

Theorem 6. *Let δ be a constant with $0 < \delta < 1$. Assume that $(n + d - 1) \leq q^{d/2}(1 - \delta)$. Then each nontrivial eigenvalue λ of the adjacency operator for the graph $G_d(n, q, \alpha)$ satisfies the bound*

$$|\lambda| \leq \frac{q^d}{d}(1 - \delta) \leq |P_d|(1 - \delta) = \lambda_{triv}(1 - \delta).$$

In particular, the graph $G_d(n, q, \alpha)$ is an expander graph.

Remarks. In our construction of the Cayley graph $G_d(n, q, \alpha)$, we took the subset E_d to be the set of all monic primary polynomials of degree d. It is also natural to take the subset to the set of all monic irreducible polynomials of degree d or the set of all monic irreducible polynomials whose degree divides d. The resulting graph would have similar qualities asymptotically. However, our choice of the subset in this paper makes the proofs simpler and cleaner with the results slightly better.

2. The number of connected components

Our key technical tool is the following Weil bound for character sums: see Theorem 2.1 in [5].

Lemma 7. *Let $\chi : \Gamma_f \longrightarrow \mathbb{C}^*$ be a nontrivial character. Then we have the estimate*

$$|\sum_{g \in P_d} \Lambda(g)\chi(g(\alpha))| \leq (n-1)\sqrt{q^d},$$

where $\Lambda(g)$ is the von Mangoldt function and it is equal to the degree of the unique prime factor in g.

Theorem 8. *If $n < q^{d/2} + 1$, then $G_d(n, q, \alpha)$ is connected.*

Proof. If the graph $G_d(n, q, \alpha)$ is not connected, then E_d generates a proper subgroup H of Γ_f. Let

$$\chi : \Gamma_f \to \Gamma_f / H \to \mathbb{C}^*$$

be a nontrivial character of Γ_f, trivial on H. Then by the Weil bound in Lemma 7,

$$q^d = \sum_{g \in P_d} \Lambda(g) = |\sum_{g \in P_d} \Lambda(g)\chi(g(\alpha))| \leq (n-1)\sqrt{q^d}.$$

It follows that $n \geq q^{d/2} + 1$. □

The next result shows that the condition $n < q^{d/2}+1$ in the above theorem is not too far from being sharp.

Theorem 9. *If $\ell > 1$ is a divisor of $(q^n - 1)$ such that $n \geq 2d + 2(|P_d| + 1)\log_q \ell$, then there is at least one $\alpha \in \mathbb{F}_{q^n}$ of degree n over $\mathbb{F}_q$ such that the number of connected components of the graph $G_d(n, q, \alpha)$ is divisible by ℓ.*

Proof. Let π_n denote the number of monic irreducible polynomials of degree n in $\mathbb{F}_q[x]$. It is easy to check that

$$|\pi_n - \frac{q^n}{n}| \leq \frac{1}{n} \sum_{k|n, k \leq n/2} q^k \leq \frac{2}{n} q^{n/2}.$$

The number of degree n elements in $\mathbb{F}_{q^n}$ is $n\pi_n$. The number of elements in $\mathbb{F}_{q^n}$ which are in a proper subfield of $\mathbb{F}_{q^n}$ containing $\mathbb{F}_q$ is

$$(\sum_{\deg(\alpha)<n} 1) = |n\pi_n - q^n| \leq 2q^{n/2}.$$

Let H be the subgroup generated by $g(\alpha)$ for $g \in P_d$. It is clear that the number of connected components of the graph $G_d(n, q, \alpha)$ is equal to the index $[\mathbb{F}_{q^n}^* : H]$.

For a divisor $\ell > 1$ of $q^n - 1$, let H_ℓ denote the unique subgroup of index ℓ in the cyclic group $\mathbb{F}_{q^n}^*$. The group H_ℓ consists of ℓ-th powers of elements in $\mathbb{F}_{q^n}^*$. Let I_d denote the set of monic irreducible polynomials g in $\mathbb{F}_q[x]$ such that $\deg(g)$ divides d. Every element of P_d is an integral power of an element in I_d. Furthermore, $|I_d| = |P_d|$. If α is a degree n element in $\mathbb{F}_{q^n}$ such that $g(\alpha) \in H_\ell$ for all $g \in I_d$, then H is a subgroup of H_ℓ and thus the number of connected components of $G_d(n, q, \alpha)$ is

$$[\mathbb{F}_{q^n}^* : H] = [\mathbb{F}_{q^n}^* : H_\ell][H_\ell : H] = \ell[H_\ell : H]$$

which is divisible by ℓ. Let

$$N_\ell = |\{\alpha \in \mathbb{F}_{q^n} : \deg(\alpha) = n,\ g(\alpha) \in H_\ell\ \forall g \in I_d\}|.$$

To prove the theorem, it is enough to prove that $N_\ell > 0$. A standard character sum argument shows that

$$\begin{aligned} \ell^{|I_d|} N_\ell &= \sum_{\deg(\alpha)=n} \prod_{g \in I_d} \sum_{\chi_g^\ell = 1} \chi_g(g(\alpha)) \\ &= \sum_{\chi_g^\ell=1, g \in I_d} \sum_{\deg(\alpha)=n} \prod_{g \in I_d} \chi_g(g(\alpha)), \end{aligned}$$

where χ_g denotes a character of $\mathbb{F}_{q^n}^*$. In the case that $\chi_g = 1$ for all $g \in I_d$, the inner sum is the number $n\pi_n$ of degree n elements in $\mathbb{F}_{q^n}$. In all other $(\ell^{|I_d|} - 1)$

cases, there is at least one $g \in I_d$ such that χ_g is a nontrivial character. In such a case, the standard Weil character sum bound (see Corollary 2.3 in [5]) implies

$$\begin{aligned}
|\sum_{\deg(\alpha)=n} \prod_{g \in I_d} \chi_g(g(\alpha))| &= |\sum_{\alpha \in \mathbb{F}_{q^n}} \prod_{g \in I_d} \chi_g(g(\alpha)) - \sum_{\deg(\alpha)<n} \prod_{g \in I_d} \chi_g(g(\alpha))| \\
&\leq ((\sum_{g \in I_d} \deg(g)) - 1) q^{n/2} + \sum_{\deg(\alpha)<n} 1 \\
&\leq (q^d - 1) q^{n/2} + 2q^{n/2} \\
&= (q^d + 1) q^{n/2},
\end{aligned}$$

where we used the fact that $\sum_{g \in I_d} \deg(g)$ is the number of elements in $\mathbb{F}_{q^d}$ of degree dividing d and thus

$$\sum_{g \in I_d} \deg(g) = q^d.$$

Putting these together, we deduce that

$$\begin{aligned}
\ell^{|I_d|} N_\ell &\geq n\pi_n - (\ell^{|I_d|} - 1)(q^d + 1) q^{n/2} \\
&\geq q^n - 2q^{n/2} - (\ell^{|I_d|} - 1)(q^d + 1) q^{n/2} \\
&\geq q^n - \ell^{|I_d|}(q^d + 1) q^{n/2} \\
&> q^{\frac{n}{2}+d}(q^{\frac{n}{2}-d} - \ell^{|I_d|+1}).
\end{aligned}$$

Solving the inequality

$$q^{\frac{n}{2}-d} \geq \ell^{|I_d|+1},$$

one obtains the condition

$$n \geq 2d + 2(|I_d| + 1) \log_q \ell.$$

Since $|I_d| = |P_d|$, the theorem is proved. □

In the case $d = 1$, we have $|I_1| = |P_1| = q$. This gives the following result.

Corollary 10. *If $\ell > 1$ is a divisor of the integer $q^n - 1$ such that $n \geq 2 + 2(q + 1) \log_q \ell$, then there is at least one degree n element α in $\mathbb{F}_{q^n}^*$ such that the graph $G_1(n, q, \alpha)$ is not connected.*

The above theorem shows that the graph $G_d(n,q,\alpha)$ is not always connected. It depends very much on the choice of α. An interesting question is to find a fast algorithm, with running time bounded by a polynomial in $(n\log q)^{O(1)}$, to compute the number of connected components. In this direction, we have the following conditional result.

Theorem 11. *Assume that the factorization of q^n-1 is given. Then one can compute the number of connected components of $G_d(n,q,\alpha)$ in time $(n\log q)^{O(1)}$.*

Proof. We may assume that $n \geq q^{d/2}+1$; otherwise $G_d(n,q,\alpha)$ is already connected. Let

$$q^n-1 = p_1^{k_1}\cdots p_s^{k_s}, H_i = \{\beta^{p_i} | \beta \in \mathbb{F}_{q^n}^*\}.$$

The H_i's are the maximal subgroups of $\mathbb{F}_{q^n}^*$. The graph $G_d(n,q,\alpha)$ is disconnected if and only if the subgroup $H =< g(\alpha) | g \in P_d >$ is contained in H_i for some i. This is true if and only if

$$g(\alpha)^{(q^n-1)/p_i} = 1, \forall g \in P_d.$$

The elements of P_d can be listed in time $q^d(n\log q)^{O(1)}$. Note that

$$\max\{s, (k_1+\cdots+k_s), q^d\} \leq n^2\log q.$$

It follows that one can check whether there is $1 \leq i \leq s$ such that $H \subseteq H_i$ in time

$$sq^d(n\log q)^{O(1)} = (n\log q)^{O(1)}.$$

If $H \not\subseteq H_i$ for $1 \leq i \leq s$, then $H = \Gamma_f$ and the graph is connected. Otherwise, we can assume that $H \subseteq H_i$ for some given i.

The group H_i is cyclic of order

$$\frac{q^n-1}{p_i} = p_1^{k_1}\cdots p_i^{k_i-1}\cdots p_s^{k_s}.$$

Its maximal subgroups are $H_{ij} = \{\beta^{p_ip_j} | \beta \in \Gamma_f\} = \{\beta^{p_j} | \beta \in H_i\}$, where $p_ip_j | (q^n-1)$. Similarly we have $H \subseteq H_{ij}$ for some j if and only if

$$g(\alpha)^{\frac{q^n-1}{p_ip_j}} = 1, \forall g \in P_d.$$

Again, we can check whether there is $1 \leq j \leq s$ such that $H \subseteq H_{ij}$ in time

$$sq^d(n\log q)^{O(1)} = (n\log q)^{O(1)}.$$

Continuing in this fashion, eventually one finds that $H = H_{i_1 i_2 \cdots i_u}$, and thus the number of connected components is $[\Gamma_f : H] = p_{i_1} \cdots p_{i_u}$. The total time needed is bounded by

$$(k_1 + \cdots + k_s)q^d(n\log q)^{O(1)} = (n\log q)^{O(1)}. \qquad \square$$

Corollary 12. *The number of connected components of $G_d(n, q, \alpha)$, which is the index $[\Gamma_f : H]$, can be computed in time $O(q^{n/4})$.*

Proof. By the well known LLL lattice factorization algorithm, $q^n - 1$ can be factored in time $O(q^{n/4})$. $\square$

Corollary 13. *If n is even, the number of connected components of $G_d(n, q, \alpha)$ can be computed in time $O(q^{n/8})$.*

Proof. $q^n - 1 = (q^{n/2} - 1)(q^{n/2} + 1)$ can be factored in time $O(q^{n/8})$. $\square$

Let $N_d(n, q, \alpha)$ denote the number of connected components of the graph $G_d(n, q, \alpha)$. An interesting problem is to give a good general upper bound for $N_d(n, q, \alpha)$, which is uniform in α. In this direction, we have the following simple crude upper bound.

Theorem 14.

$$N_d(n, q, \alpha) \leq \frac{q^n - 1}{\binom{|P_d| + \lceil \frac{n}{d} \rceil}{\lceil \frac{n}{d} \rceil - 1}}.$$

Proof. Let H be the subgroup generated by $\{g(\alpha) | g \in P_d\}$. Since α has degree n, the unique factorization of polynomials implies that the elements

$$g_1(\alpha) \cdots g_k(\alpha), 0 \leq k \leq \lceil \frac{n}{d} \rceil - 1, \{g_1, \cdots, g_k\} \subset P_d$$

are elements of H. This proves

$$|H| \geq \binom{|P_d| + \lceil \frac{n}{d} \rceil}{\lceil \frac{n}{d} \rceil - 1}.$$

It follows that

$$N_d(n, q, \alpha) = [\Gamma_f : H] = \frac{q^n - 1}{|H|} \leq \frac{q^n - 1}{\binom{|P_d| + \lceil \frac{n}{d} \rceil}{\lceil \frac{n}{d} \rceil - 1}}.$$

The theorem is proved. $\square$

3. The diameter

The diameter of $G_d(n, q, \alpha)$ is the minimal integer D (or ∞ if it does not exist) such that every element in Γ_f can be written as a product of at most D elements in E_d.

Theorem 15. *Assume that $n < q^{d/2} + 1$. The diameter D of $G_d(n, q, \alpha)$ satisfies the inequality*

$$D \leq 2\frac{n}{d} + 1 + \frac{4\frac{n}{d}\log(n-1)}{d\log q - 2\log(n-1)}.$$

Proof. Let $\widehat{\Gamma_f}$ be the character group of the multiplicative group $\Gamma_f = \mathbb{F}_{q^n}^*$, which is the set of homomorphisms from Γ_f to $\mathbb{C}^*$. For integer $k > 0$ and $\beta \in \Gamma_f$, let $N_k(\beta)$ be the number of solutions of the equation

$$\beta = g_1(\alpha)g_2(\alpha)\cdots g_k(\alpha), g_i \in P_d.$$

It is clear that

$$N_k(\beta) = \frac{1}{q^n - 1}\sum_{g_1,\cdots,g_k \in P_d}\sum_{\chi \in \widehat{\Gamma_f}} \chi(\frac{g_1(\alpha)\cdots g_k(\alpha)}{\beta}).$$

To show that the diameter D is bounded by k, it is enough to show that $N_k(\beta) > 0$ for all $\beta \in \Gamma_f$. For our purpose, it is simpler to work with the following weighted sum

$$M_k(\beta) = \frac{1}{q^n - 1}\sum_{g_1,\cdots,g_k \in P_d} \Lambda(g_1)\cdots\Lambda(g_k)\sum_{\chi \in \widehat{\Gamma_f}} \chi(\frac{g_1(\alpha)\cdots g_k(\alpha)}{\beta}).$$

Note that $N_k(\beta) > 0$ if and only if $M_k(\beta) > 0$. Now, separating the trivial character, we obtain

$$\begin{aligned} M_k(\beta) &= \frac{q^{kd}}{q^n - 1} + \frac{1}{q^n - 1}\sum_{g_1,\cdots,g_k \in P_d} \Lambda(g_1)\cdots\Lambda(g_k)\sum_{\chi \neq 1} \chi(\frac{g_1(\alpha)\cdots g_k(\alpha)}{\beta}) \\ &= \frac{q^{kd}}{q^n - 1} + \frac{1}{q^n - 1}\sum_{\chi \neq 1} \chi^{-1}(\beta)(\sum_{g \in P_d} \Lambda(g)\chi(g(\alpha)))^k. \end{aligned}$$

Applying the Weil bound in Lemma 7, we deduce that

$$|M_k(\beta) - \frac{q^{kd}}{q^n - 1}| < (n-1)^k \sqrt{q^{dk}}.$$

In order for $M_k(\beta) > 0$ for all β, it suffices to have the inequality

$$q^{kd} \geq q^n (n-1)^k q^{kd/2},$$

that is,

$$q^{kd-2n} \geq (n-1)^{2k}.$$

This is satisfied if

$$k \geq \frac{2n}{d - 2\log_q(n-1)} = 2\frac{n}{d} + \frac{4\frac{n}{d}\log(n-1)}{d\log q - 2\log(n-1)}.$$

The theorem is proved. □

For a proper divisor d of n, we now make some comparisons between Chung's graph $G_1(\frac{n}{d}, q^d, \beta)$ and our more general construction $G_d(n, q, \alpha)$, where β is a root of an irreducible polynomial of degree n/d in $\mathbb{F}_{q^d}[x]$ and α is a root of an irreducible polynomial of degree n in $\mathbb{F}_q[x]$. It is clear that both graphs have $q^n - 1$ vertices. Assume that $n < q^{d/2} + 1$. In this case, both $G_1(\frac{n}{d}, q^d, \beta)$ and $G_d(n, q, \alpha)$ are connected, and their diameter bounds

$$D_1 \leq 2\frac{n}{d} + 1 + \frac{4\frac{n}{d}\log(\frac{n}{d} - 1)}{d\log q - 2\log(\frac{n}{d} - 1)}, \; D_2 \leq 2\frac{n}{d} + 1 + \frac{4\frac{n}{d}\log(n-1)}{d\log q - 2\log(n-1)}$$

are comparable. But $G_1(\frac{n}{d}, q^d, \alpha)$ is q^d-regular and $G_d(n, q, \alpha)$ is $|P_d|$-regular, where $|P_d| \sim \frac{q^d}{d} < q^d$. Thus, $G_d(n, q, \alpha)$ can be significantly better than $G_1(\frac{n}{d}, q^d, \alpha)$ if $n < q^{d/2} + 1$, since $G_d(n, q, \alpha)$ has far fewer edges.

Corollary 16. *If $q^d > (n-1)^{4\frac{n}{d}+2}$, then $D \leq 2\frac{n}{d} + 1$.*

If q is sufficiently large, it may be possible to improve the above diameter bound to $D \leq \frac{n}{d} + 2$. This is indeed the case for $d = 1$, as shown by Katz [4] and Cohen [2].

A computational question is to ask for a fast algorithm, with running time bounded by $O(n \log q)^{O(1)}$, to compute the diameter D of the graph $G_d(n, q, \alpha)$. This is expected to be a very difficult problem. Even assuming the factorization of $q^n - 1$, we still do not know a fast algorithm to compute the diameter. We believe that computing the diameter is related to the discrete logarithm problem and the subset sum problem, both of which are difficult problems used in cryptography.

4. Expander graphs

In this section, we show that our graph $G_d(n, q, \alpha)$ has good expanding properties. The adjacency matrix $M = (m_{\beta_1\beta_2})$ is a $(q^n-1)\times(q^n-1)$ matrix, where the entry $m_{\beta_1\beta_2} = 1$ if $\beta_1 \to \beta_2$ is an edge and it is zero otherwise. The adjacency operator M acts on the $(q^n - 1)$-dimensional complex vector space $\mathbb{C}^{\Gamma_f}$ of functions on Γ_f. If $h(x)$ is a complex function on Γ_f, then

$$M(h)(x) = \sum_{x\to y} h(y) = \sum_{g\in P_d} h(xg(\alpha)),$$

where y runs over all elements of Γ_f such that $x \to y$ is an edge of $G_d(n, q, \alpha)$. If $h(x) = \chi(x)$ is a multiplicative character of Γ_f, then one checks that

$$M(\chi)(x) = \sum_{g\in P_d} \chi(xg(\alpha)) = \lambda_d(\chi)\chi(x),$$

where

$$\lambda_d(\chi) = \sum_{g\in P_d} \chi(g(\alpha)).$$

This shows that each character χ is an eigenvector of the operator M. By Artin's lemma, the set of characters on Γ_f is $\mathbb{C}$-linearly independent. Since the number of characters is equal to $q^n - 1$, it follows that $\mathbb{C}^{\Gamma_f}$ has a basis consisting of the eigenvectors χ of M, where χ runs through all characters of Γ_f. If χ is a character which is trivial on the subgroup generated by $H =< g(\alpha)|g \in P_d >$ of Γ_f, then the eigenvalue

$$\lambda_d(\chi) = \sum_{g\in P_d} 1 = |P_d|$$

which is the trivial eigenvalue $\lambda_{triv} = |P_d|$. If χ is a character which is non-trivial on H, its eigenvalue is called a nontrivial eigenvalue which satisfies the bound

$$
\begin{aligned}
|\lambda_d(\chi)| &= |\sum_{g\in P_d}\chi(g(\alpha))| \\
&= |\frac{1}{d}\sum_{g\in P_d}\Lambda(g)\chi(g(\alpha)) + \sum_{g\in P_d,\Lambda(g)<d}(1-\frac{\Lambda(g)}{d})\chi(g(\alpha))| \\
&\leq \frac{n-1}{d}q^{d/2} + \sum_{g\in P_d,\Lambda(g)<d}(1-\frac{\Lambda(g)}{d}) \\
&\leq \frac{n-1}{d}q^{d/2} + q^{d/2} \leq \frac{n+d-1}{d}q^{d/2}.
\end{aligned}
$$

Since

$$\frac{q^d}{d} = \sum_{g\in P_d}\frac{\Lambda(g)}{d} \leq \sum_{g\in P_d} 1 = |P_d|,$$

we deduce

Theorem 17. *Let δ be a constant with $0<\delta<1$. Assume that $(n+d-1)\leq q^{d/2}(1-\delta)$. Then each nontrivial eigenvalue λ of the adjacency operator M for the graph $G_d(n,q,\alpha)$ satisfies the bound*

$$|\lambda| \leq \frac{q^d}{d}(1-\delta) \leq |P_d|(1-\delta) = \lambda_{triv}(1-\delta).$$

In particular, the graph $G_d(n,q,\alpha)$ is an expander graph.

Note that the number of connected components of $G_d(n,q,\alpha)$ is equal to the multiplicity of the trivial eigenvalue $|P_d|$ of the adjacency matrix M. If one uses the matrix M and linear algebra directly to compute the number of connected components, then the running time will be $O(q^n)^{O(1)}$, which is fully exponential in terms of $n\log q$. This trivial algorithm is far slower than the conditional result in Theorem 5.

Finally, we explain that it is best to view our graph $G_d(n,q,\alpha)$ as a weighted graph. For this purpose, let $G_d^*(n,q,\alpha)$ be the weighted graph with the same vertices and edges as $G_d(n,q,\alpha)$. Given an edge $\beta_1\to\beta_2$ in $G_d^*(n,q,\alpha)$, we define the weight of the edge $\beta_1\to\beta_2$ to be $\Lambda(\beta_2/\beta_1)=\Lambda(g)$, where $\beta_2/\beta_1=g(\alpha)$ for a unique monic primary polynomial $g\in P_d$. The weighted adjacency matrix $M^*=(m_{\beta_1\beta_2})$ is a $(q^n-1)\times(q^n-1)$ matrix, where the entry $m_{\beta_1\beta_2} = \Lambda(\beta_2/\beta_1)$ if $\beta_1 \to \beta_2$ is an edge and it is zero otherwise. The adjacency operator M^* acts on the $(q^n - 1)$-dimensional complex vector space $\mathbb{C}^{\Gamma_f}$ of functions on Γ_f. If $h(x)$ is a complex function on Γ_f, then

$$M^*(h)(x) = \sum_{x \to y} \Lambda(\frac{y}{x})h(y) = \sum_{g \in P_d} \Lambda(g)h(xg(\alpha)),$$

where y runs over all elements of Γ_f such that $x \to y$ is an edge of $G_d^*(n, q, \alpha)$. If $h(x) = \chi(x)$ is a multiplicative character of Γ_f, then one checks that

$$M^*(\chi)(x) = \sum_{g \in P_d} \Lambda(g)\chi(xg(\alpha)) = S_d(\chi)\chi(x),$$

where

$$S_d(\chi) = \sum_{g \in P_d} \Lambda(g)\chi(g(\alpha)).$$

This shows that each character χ is an eigenvector of the operator M^*. If χ is a character which is trivial on the subgroup generated by $H =< g(\alpha)|g \in P_d >$ of Γ_f, then the eigenvalue

$$S_d(\chi) = \sum_{g \in P_d} \Lambda(g) = q^d$$

which is the trivial eigenvalue $\lambda_{triv} = q^d$. If χ is a character which is nontrivial on H, its eigenvalue is called a nontrivial eigenvalue which satisfies the bound

$$|S_d(\chi)| = |\sum_{g \in P_d} \Lambda(g)\chi(g(\alpha))| \leq (n-1)\sqrt{q^d}.$$

We obtain

Theorem 18. *Let δ be a constant with $0 < \delta < 1$. Assume that $(n-1) \leq q^{d/2}(1-\delta)$. Then each nontrivial eigenvalue λ of the adjacency operator M^* for the weighted graph $G_d^*(n, q, \alpha)$ satisfies the bound*

$$|\lambda| \leq \lambda_{triv}(1-\delta).$$

In particular, the weighted graph $G_d^*(n, q, \alpha)$ is an expander graph.

The condition $(n-1) \leq q^{d/2}(1-\delta)$ in this weighted theorem is weaker and simpler than the condition $(n+d-1) \leq q^{d/2}(1-\delta)$ in the previous unweighted theorem.

[1] F. R. Chung, Diameter and eigenvalues, J. Amer. Math. Soc., 2 (1989), 187-196.

[2] S. D. Cohen, Polynomial factorisation and an application to regular directed graphs, Finite Fields Appl., 4 (1998), 316-346.

[3] S. Hoory, N. Linial and A. Wigderson, Expander graphs and their applications, Bull. Amer. Math. Soc., (New Series), 43(2006), 439-561.

[4] N. Katz, Factoring polynomials in finite fields: an application of Lang-Weil to a problem in graph theory, Math. Ann., 286 (1990), no. 4, 625-637.

[5] D. Wan, Generators and irreducible polynomials over finite fields, Math. Comp., 66(1997), no. 219, 1195-1212.