The chromatic number of random Cayley graphs

Noga Alon *

Abstract

We consider the typical behaviour of the chromatic number of a random Cayley graph of a given group of order n with respect to a randomly chosen set of size $k \leq n/2$. This behaviour depends on the group: for some groups it is typically 2 for all $k < 0.99 \log_2 n$, whereas for some other groups it grows whenever k grows. The results obtained include a proof that for any large prime p, and any $1 \leq k \leq 0.99 \log_3 p$, the chromatic number of the Cayley graph of Z_p with respect to a uniform random set of k generators is, asymptotically almost surely, precisely 3. This strengthens a recent result of Czerwiński. On the other hand, for $k > \log p$, the chromatic number is typically at least $\Omega(\sqrt{k/\log p})$ and for $k = \Theta(p)$ it is $\Theta(\frac{p}{\log p})$.

For some non-abelian groups (like $SL_2(Z_q)$), the chromatic number is, asymptotically almost surely, at least $k^{\Omega(1)}$ for every k, whereas for elementary abelian 2-groups of order $n = 2^t$ and any k satisfying $1.001t \leq k \leq 2.999t$ the chromatic number is, asymptotically almost surely, precisely 4. Despite these discrepancies between different groups, it seems plausible to conjecture that for any group of order n and any $k \leq n/2$, the typical chromatic number of the corresponding Cayley graph cannot differ from k by more than a poly-logarithmic factor in n.

1 Introduction

Let B be a finite group of order n. For an integer $k \leq n/2$, let S be a random subset of B obtained by choosing, randomly, uniformly and independently (with repetitions), k elements of G, and by letting S be the set of these elements and their inverses, without the identity. Thus S is a set of cardinality at most $2k$, and is typically of cardinality at least $k - O(k^2/n)$. In this paper we consider the behaviour of the chromatic number of the Cayley graph of B with respect to S, that is, the graph whose vertices are all members of B where b_1 and b_2 are adjacent if $b_1 \cdot b_2^{-1} \in S$. We denote this graph by (B, k), and its chromatic number by $\chi(B, k)$.

One motivation for studying this problem is the constructions in [6] in which random self complementary Cayley graphs of high chromatic number are used in the investigation of a problem in Information Theory, providing graphs with a big gap between their chromatic number and their so-called Witsenhausen rate-see [6] for more details. Another motivation is the fact that many of the

*Sackler School of Mathematics and Blavatnik School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel and Institute for Advanced Study, Princeton, New Jersey, 08540, USA. Email: nogaa@tau.ac.il. Research supported in part by an ERC Advanced grant, by a USA-Israeli BSF grant and by NSF grant No. DMS-0835373.

known constructions of expanders, like the ones in [5], [18], [19] are Cayley graphs, the fact that random Cayley graphs with logarithmic degrees over any group are typically expanders [8], and over some groups even a bounded degree suffices [9], and the fact that graphs with strong expansion properties have high chromatic numbers. Yet another reason is the study of an extremal problem of Green regarding sumsets in finite fields, whose investigation leads to the question of estimating the typical behaviour of the chromatic number of random Cayley sum graphs of Z_p. See [13], [1] for more details.

Our results are asymptotic and we are interested in the case of large n, where k may grow with n or stay constant. As usual, we will say that a property holds asymptotically almost surely (a.a.s., for short), if the probability it holds tends to 1 as n tends to infinity. The problem of determining the typical asymptotic behaviour of $\chi(B,k)$ for a general given group B of order n and general $k \leq n/2$ appears to be very difficult, but we do obtain several nontrivial estimates for general groups, as well as more accurate estimates for specific groups.

The rest of the paper is organized as follows. In the next section we consider general groups, cyclic groups are considered in Section 3 and abelian ones in Section 4. The final Section 5 contains several open problems. Throughout the paper all logarithms are in base 2, unless otherwise specified, and we omit floor and ceiling signs whenever these are not crucial. We generally make no serious attempts to optimize the absolute constants in (most of) our estimates.

2 General groups

Note that (B,k) is regular of degree at most $2k$, and hence always $\chi(B,k) \leq 2k+1$.

Theorem 2.1 *For any group B of order n and any $k \leq n/2$, the chromatic number $\chi(B,k)$ satisfies, a.a.s, the following bounds.*
(i) $\chi(B,k) \leq O(k/\log k)$.
(ii) $\chi(B,k) \geq \Omega((k/\log n)^{1/2})$.
(iii) $\chi(B,k) \geq \Omega(\frac{k^2}{n \log^2 n})$.

In order to prove the theorem we need several lemmas. The first two supply upper bounds for the chromatic number of sparse or pseudo-random graphs.

Lemma 2.2 ([4]) *The chromatic number of any graph with maximum degree d in which every neighborhood of a vertex spans at most d^2/f edges, where $f < d^2$, is $O(d/\log f)$.*

Lemma 2.3 ([3]) *The chromatic number of any d-regular graph with all nontrivial eigenvalues bounded in absolute value by λ is at most*

$$O(\frac{d-\lambda}{\log(\frac{d-\lambda}{\lambda+1}+1)}).$$

The following lemma is proved in [8]. see also [16], [17], [11] for subsequent alternative proofs, providing somewhat sharper estimates

Lemma 2.4 ([8]) *Let G be a random Cayley graph of a group of order n with a random set S of k generators. Then, a.a.s., every nontrivial eigenvalues of G is, in absolute value, at most*

$$\lambda = O(\sqrt{k}\sqrt{\log n}).$$

We will also apply the following well known result of Hoffman.

Lemma 2.5 ([15]) *Let G be a d-regular graph on n vertices in which the smallest (that is, the most negative) eigenvalue is λ_n. Then the maximum size of an independent set in G is at most $\frac{-n\lambda_n}{d-\lambda_n}$ and hence its chromatic number is at least $-\frac{d-\lambda_n}{\lambda_n}$.*

Call a subset D of cardinality $|D| = t$ in a group B a *subset with many quotients* if the number of distinct elements of the form $d(d')^{-1}$ with $d, d' \in D$ is at least $d^2/5$. The following simple lemma asserts that any set C in a group contains a subset with many quotients of size at least $\lfloor\sqrt{|C|}\rfloor$. A version of this lemma for abelian groups appears in [6].

Lemma 2.6 *Any set C of cardinality $r > r_0$ in a group contains a subset D with many quotients of cardinality at least $\lfloor\sqrt{r}\rfloor$.*

Proof. Let $C = \{c_1, c_2, \dots, c_r\}$, and let $D = \{\ c_{i_1}, c_{i_2}, \dots, c_{i_t}\ \} = \{d_1, d_2, \dots, d_t\}$, with $d_j = c_{i_j}$ and $1 \le i_1 < i_2 < \cdots < i_t \le r$, be a random subset of t elements of C, where t will be chosen later. As the group may contain elements of order 2 and thus it may be that $d(d')^{-1} = d'd^{-1}$ for some distinct $d, d' \in D$, we will consider only quotients of the form $d_i d_j^{-1}$ with $i < j$. There are $\binom{t}{2}$ such quotients, but some of them may be equal.

Let $n_3(D)$ denote the number of ordered triples (d_i, d_j, d_k) of elements of D with $i < j < k$ so that $d_i(d_j)^{-1} = d_j(d_k)^{-1}$. Let $n_4(D)$ denote the number of ordered 4-tuples (d_i, d_j, d_k, d_ℓ) of elements of D with $i < j$ and $k < \ell$, so that $d_i(d_j)^{-1} = d_k(d_\ell)^{-1}$. Similarly, let $n_3(C)$ and $n_4(C)$ be defined analogously, with respect to the larger set C.

It is not difficult to check that the number of distinct quotients of the form $d_i(d_j)^{-1}$ with $i < j$, $d_i, d_j \in D$ is at least $\binom{t}{2} - n_3(D) - \frac{1}{2}n_4(T)$. Indeed, each group element obtained $h > 1$ times as a quotient of the above form, contributes at least $h - 1$ to $n_3(D) + \frac{1}{2}n_4(D)$.

We proceed to estimate the expectation of the random variable $\binom{t}{2} - n_3(D) - \frac{1}{2}n_4(D)$. The expected values $E(n_3(D))$ and $E(n_4(D))$ satisfy

$$E(n_3(D)) = E(n_3(C))\frac{t(t-1)(t-2)}{r(r-1)(r-2)},$$

and

$$E(n_4(D)) = E(n_4(C))\frac{t(t-1)(t-2)(t-3)}{r(r-1)(r-2)(r-3)}.$$

Indeed, this follows, by linearity of expectation, from the fact that the probability that a fixed set of 3 members of C is contained in D is exactly $\frac{t(t-1)(t-2)}{r(r-1)(r-2)}$, and an analogous expression gives the probability that a fixed set of 4 members of C is a subset of D.

It is also easy to see that $n_3(C) \le \binom{r-1}{2}$, as there are that many ways to choose $c_i, c_j \in C$ with $i < j < r$, and this determines a unique group element g so that $c_i(c_j)^{-1} = c_j g^{-1}$ (which may be some c_k for $k > j$.) A similar argument implies that $n_4(C) \le \binom{r}{2}(r-2)$, as there are $\binom{r}{2}(r-2)$ ways to choose three distinct elements of C, (c_i, c_j, c_k) with $i < j$ and this determines a unique group element g with $c_i c_j^{-1} = c_k g^{-1}$. It thus follows, by linearity of expectation, that the expectation of the difference $\binom{t}{2} - n_3(D) - \frac{1}{2} n_4(D)$ is a least

$$\binom{t}{2} - \frac{(r-1)(r-2)}{2} \frac{t(t-1)(t-2)}{r(r-1)(r-2)} - \frac{r(r-1)(r-2)}{4} \frac{t(t-1)(t-2)(t-3)}{r(r-1)(r-2)(r-3)}$$

$$= \frac{t(t-1)}{2}[1 - \frac{t-2}{r} - \frac{(t-2)(t-3)}{2(r-3)}].$$

Taking $t = \lfloor \sqrt{r} \rfloor$ we conclude that for large t, the expected number of distinct quotients of elements of D, which is at least the expectation of $\binom{t}{2} - n_3(D) - \frac{1}{2} n_4(D)$, is a least $(\frac{1}{4} - o(1))t^2 > \frac{1}{5} t^2$. Thus there exists a T with many quotients, as needed. $\square$

Proof of Theorem 2.1:

(i) For, say, $k < n^{1/4}$ we apply Lemma 2.2. The random graph we consider is a Cayley graph of the group B with respect to a set S of cardinality at most $2k$. This is an $|S|$-regular graph and the number of edges in a neighborhood of each of its vertices is at most the number of ordered triples (s_1, s_2, s_3) where $s_i \in S$ for all i and $s_1 \cdot s_2 \cdot s_3 = 1$, where 1 is the identity of B. The number of such triples in which one of the elements s_i is equal to another one s_j or to its inverse is clearly $O(k)$. The expected number of such triples in which no element s_i is equal to another one or to its inverse is at most $O(n^2(\frac{k}{n})^3) = O(\frac{k^3}{n}) = O(n^{-1/4})$, as the number of ordered triples (x_1, x_2, x_3) of elements of B whose product is 1 is at most n^2, and the probability that all members of such a triple belong to the random set S is $O((\frac{k}{n})^3)$. It thus follows, by Markov's inequality, that a.a.s. the number of edges in a neighborhood of a vertex is at most $O(k)$, and the required $O(k/\log k)$ upper bound for the chromatic number follows from Lemma 2.2

For $k \ge n^{1/4}$, we have $\log k = \Theta(\log n)$. Here we use the result of [8] quoted as Lemma 2.4 above. As we consider now the range $k \ge n^{1/4}$, $\lambda < k^{3/4}$. We can now apply Lemma 2.3. In our case, a.a.s, $d = \Theta(k)$ and $\lambda < k^{3/4}$, providing the required bound and completing the proof of (i).

(ii) The required assertion follows immediately from lemma 2.4 and Lemma 2.5.

(iii) Define $t = \frac{5n \ln n}{k}$ and note that as $k \le n/2$ this number is at least $10 \ln n$. We claim that a.a.s. the random Cayley graph (B, k) contains no independent set D of size t which forms a set with many quotients. Indeed, the probability that a fixed such set is independent is at most the probability that all k random choices of the elements in our generating set lie outside the set $D \cdot D^{-1}$, whose cardinality is at least $\frac{t^2}{5}$. This probability is at most $(1 - \frac{t^2}{5n})^k \le e^{-t^2 k/(5n)}$. As the number of potential sets D is $\binom{n}{t} \le \frac{n^t}{t!}$ we conclude that the probability that there exists such a D is at most

$$\frac{1}{t!} n^t \cdot e^{-t^2 k/(5n)} \le \frac{1}{t!} [n e^{-tk/5n}]^t = \frac{1}{t!}.$$

As $t > 10 \ln n$ this probability is negligible, proving the claim. By Lemma 2.6 this implies that a.a.s our graph contains no independent set of size $t^2 = \frac{25n^2 \ln^2 n}{k^2}$ supplying the desired lower bound for the chromatic number. $\Box$

We conclude this section by observing that for some groups B, the typical chromatic number $\chi(B,k)$ grows with k even if k is very small as a function of the size n of the group. (As we show later, this is not the case if B is abelian).

Proposition 2.7 *There exists an absolute constant $\delta > 0$ so that if B is the group $SL_2(Z_q)$ then for any k, $\chi(B,k) \geq \Omega(k^{\delta})$ a.a.s.*

The proof is an immediate consequence of Lemma 2.5 and the following result of Bourgain and Gamburd.

Lemma 2.8 ([9]) *There exists an absolute constant $\delta > 0$ so that if B is the group $SL_2(Z_q)$ then for any k, every nontrivial eigenvalue of the random Cayley graph (B,k) is, in absolute value, at most $k^{1-\delta}$ a.a.s.*

3 Cyclic groups

3.1 Groups of prime order

A recent result of Czerwiński [10] implies that for any prime p and for $k \leq (\frac{\log p}{\log \log p})^{1/2}$, $\chi(Z_p, k) = 3$ a.a.s. It turns out that this holds for a wider range of k including all $k \leq (1 - o(1)) \log_3 p$. This is proved in Theorem 3.2 below. Note that by Theorem 2.1, part (ii), once $k > C \log p$ for some large constant C, the chromatic number exceeds 3 (and is at least $\Omega(\sqrt{C})$) a.a.s.

The main tool in the proof of Theorem 3.2 is the following lemma.

Lemma 3.1 *Let p be an odd prime, let δ, μ be positive reals satisfying $1 > \delta > 2\mu > 0$ and let I be a cyclic interval in Z_p of size $|I| = \delta p$. Let $A \subset Z_p$ be an arbitrary subset of Z_p, and let x be a uniformly chosen random element of Z_p. Define $A' = \{a \in A \ : \ xa \in I\}$. Then the probability that the size of A' is smaller than $(\delta - 2\mu)|A|$ satisfies*

$$Pr(|A'| < (\delta - 2\mu)|A|) \leq \lceil \frac{1}{\mu} \rceil \frac{(\delta - \mu)(1 - \delta + \mu)|A|}{\mu^2 |A|^2} < \frac{2\delta(1 - \delta + \mu)}{\mu^3 |A|}.$$

Proof. We apply the second moment method, a similar application appears in [7] and in [2]. Put $r = \lceil \frac{1}{\mu} \rceil$ and let $L = \{J_1, J_2, \ldots, J_r\}$ be a family of cyclic intervals in Z_p, each of size $(\delta - \mu)p$, so that any cyclic interval I of size δp fully contains at least one J_i. It is clear that such a set of intervals J_i exists, simply choose their leftmost points with (nearly) equal spacing in Z_p. Fix an interval J in L, and let y be another random uniform member of Z_p, independent of x. For each $a \in A$, put $z_a = ax + y$, and let Z_a be the indicator random variable whose value is 1 if and only if $z_a \in J$. Define also $Z = \sum_{a \in A} Z_a$. Note that each z_a is uniformly distributed in Z_p and hence the expectation of Z_a is exactly $\delta - \mu$. Moreover, crucially, for each two distinct $a, a' \in A$, the ordered pair $(z_a, z_{a'})$ is uniformly distributed in Z_p^2, implying that the random variables $\{Z_a : a \in A\}$ are pairwise independent. It follows that the expectation of $Z = \sum_{a \in A} Z_a$ is $(\delta - \mu)|A|$, and its variance is $(\delta - \mu)(1 - \delta + \mu)|A|$. By Chebyshev's Inequality, the probability that the value of Z deviates from its expectation by at least $\mu|A|$ is at most

$$\frac{\mathrm{Var}[Z]}{\mu^2|A|^2} = \frac{(\delta - \mu)(1 - \delta + \mu)|A|}{\mu^2|A|^2}.$$

Therefore, the probability that there exists an interval J in L which contains less than $(\delta - 2\mu)|A|$ elements z_a is at most r times the above bound. It follows that with probability at least

$$P = 1 - r\frac{(\delta - \mu)(1 - \delta + \mu)|A|}{\mu^2|A|^2}$$

(over the choices of x and y) every interval in L contains at least $(\delta - 2\mu)|A|$ elements z_a, and hence there is a fixed y so that for this y, as x is chosen at random, the probability that every interval in L contains at least that many numbers z_a is at least P. However, by the construction of the family L, the interval $I + y$ (whose length is δp) fully contains one of the intervals in L, and hence with the above probability it contains at least $(\delta - 2\mu)|A|$ elements of the form $z_a = ax + y$, implying that I contains that many elements ax. This completes the proof. $\square$

Theorem 3.2 *For any fixed $\epsilon > 0$, if p is a prime and $1 \le k \le (1 - \epsilon)\log_3 p$ then the chromatic number $\chi(Z_p, k)$ is, a.a.s., exactly* 3.

Proof. As the order of each nonzero member of Z_p is p, which is odd, the chromatic number is at least 3. To prove the upper bound, let $S = (x_1, x_2 \ldots, x_k)$ with $k \le (1 - \epsilon)\log_3 p$ be a sequence of random elements of Z_p, and consider the Cayley graph of Z_p with respect to $S \cup (-S)$. Let $I = \{\lceil p/3 \rceil, \ldots, \lfloor 2p/3 \rfloor\}$ be the interval consisting of (roughly) the third middle of Z_p. Note that its size is δp where $\delta = \frac{1}{3} + \Theta(1/p)$ and the $\Theta(1/p)$ term is positive for $p \equiv 2(mod\ 3)$ and is negative for $p \equiv 1(mod\ 3)$.

Claim: A.a.s. there exists an $a \in Z_p$ so that $ax_i \in I$ for all $1 \le i \le k$.

Proof of Claim: Let μ be a fixed real satisfying $\delta > 2\mu$ (its exact value will be chosen later). Put $A_0 = Z_p$ and for each i, $1 \le i \le k$, define $A_i = \{a \in A_{i-1}, ax_i \in I\}$. By Lemma 3.1, $|A_i| \ge (\delta - 2\mu)|A_{i-1}|$ with probability at least $1 - \frac{2\delta(1-\delta+\mu)}{\mu^3|A_{i-1}|}$. Therefore, with probability at least

$$1 - \frac{2\delta(1 - \delta + \mu)}{\mu^3 p}\sum_{i=0}^{k-1}\frac{1}{(\delta - 2\mu)^i} = 1 - \frac{2\delta(1 - \delta + \mu)}{\mu^3 p}[\frac{1}{(\delta - 2\mu)^k} - 1]/[\frac{1}{\delta - 2\mu} - 1] \ge 1 - \frac{2\delta}{\mu^3 p(\delta - 2\mu)^{k-1}},$$

we have $|A_i| \ge (\delta - 2\mu)^i p$ for all i. In our case $\delta \ge \frac{1}{3} - \frac{1}{3p}$. By choosing, say, $\mu = \frac{\epsilon}{20}$ this implies that with probability at least $1 - p^{-\Omega(\epsilon)}$ the set A_k is nonempty. Any $a \in A_k$ satisfies the assertion of the claim. $\square$

Having proved the claim we can now complete the proof of the theorem. Cover Z_p by three pairwise disjoint intervals I_1, I_2, I_3, each having at most $\lceil p/3 \rceil$ elements, and color any $z \in Z_p$ by the index i, $1 \leq i \leq 3$ such that in Z_p, $az \in I_i$. It is easy to check that this is a proper coloring. Indeed, if z_1 and z_2 have the same color then az_1 and az_2 lie in the same interval I_j and thus differ in Z_p by at most $\lceil p/3 \rceil - 1 = \lfloor p/3 \rfloor$. Therefore their difference cannot be of the form ax_i or $-ax_i$ for some i, as all these numbers lie in $I = \{\lceil p/3 \rceil, \ldots, \lfloor 2p/3 \rfloor\}$. $\square$

Remarks:

- The proof above works for elementary abelian p-groups $B = Z_p^m$ by a simple modification, showing that for any odd prime p, whenever k does not exceed $c \log |B|$ for an absolute positive constant c, then $\chi(B, k) = 3$ a.a.s.
- For two integers $a \geq 2b$, an (a, b)-coloring of a graph $G = (V, E)$ is a mapping $f : V \mapsto \{0, 1, \ldots, a-1\}$ so that for every edge $uv \in E$, $b \leq |f(u) - f(v)| \leq a - b$. The circular chromatic number $\chi_c(G)$ of G is the minimum possible ratio a/b so that there is an (a, b)-coloring of G. See [21] for a survey on circular coloring. It is known that for any graph G, $\lceil \chi_c(G) \rceil = \chi(G)$. The proof above can be easily modified to show that for any positive μ there is a positive $c = c(\mu)$ so that for any large prime p, if $k \leq c \log p$ then the circular chromatic number of a random Cayley graph of Z_p with respect to k randomly chosen generators is, asymptotically almost surely, at most $2 + \mu$.

By a similar reasoning, we can prove the following.

Theorem 3.3 *For all k, $1 < k \leq p/2$, $\chi(Z_p, k) \leq O(1 + \frac{k}{\log p})$ a.a.s.*

Proof. Observe, first, that the the statement for $k < 10 \log p$, say, follows from the one for $k = 10 \log p$ and that for $k \geq p^{0.1}$ the statement follows from Theorem 2.1, part (i). We thus may and will assume that $10 \log p \leq k \leq p^{0.1}$.

Put $k = g \log p$, and let $S = \{x_1, \ldots, x_k\}$ be a random sequence of elements of Z_p. Note that $10 \leq g < p^{0.1}/\log p$. Define I to be the following interval in Z_p,

$$I = \{\lceil \frac{p}{10g} \rceil, \lceil \frac{p}{10g} \rceil + 1, \ldots, p - \lceil \frac{p}{10g} \rceil\}.$$

Thus $|I| = \delta p$ where $\delta = 1 - \frac{2}{10g} + \Theta(\frac{1}{p})$. Put $\mu = \frac{1}{10g}$, then for large p,

$$\delta - 2\mu = 1 - \frac{4}{10g} + \Theta(\frac{1}{p}) \geq 1 - \frac{1}{2g}.$$

As in the previous proof, we claim that a.a.s. there exists an $a \in Z_p$ so that $ax_i \in I$ for all $1 \leq i \leq k$. To prove this claim define, as before, $A_0 = Z_p$ and $A_i = \{a \in A_{i-1}, ax_i \in I\}$. Thus $|A_i| \geq (\delta - 2\mu)^i p$ for all $1 \leq i \leq k$ with probability at least

$$1 - \frac{2\delta(1 - \delta + \mu)}{\mu^3 p} \sum_{i=0}^{k-1} \frac{1}{(\delta - 2\mu)^i} = 1 - \frac{2\delta(1 - \delta + \mu)}{\mu^3 p} [\frac{1}{(\delta - 2\mu)^k} - 1] / [\frac{1}{\delta - 2\mu} - 1] \geq 1 - \frac{2\delta}{\mu^3 p (\delta - 2\mu)^{k-1}}.$$

In our case $\mu^3 \geq p^{-0.3}$ and $(\delta - 2\mu)^{k-1} \geq p^{-2/3}$, implying that with probability at least $1 - p^{-\Omega(1)}$ A_k is nonempty, providing the existence of the claimed a.

Once we have the claimed a it is easy to define a proper coloring by $O(g)$ colors. Indeed, split Z_p into, say, $20g$ nearly equal pairwise disjoint intervals and color each $z \in Z_p$ by the index of the interval containing az. If two elements z, z' have the same color then az, az' lie in the same interval, and hence the difference between these two lies in $[-\frac{p}{20g}, \frac{p}{20g}]$ and therefore cannot be an ax_i or $-ax_i$, as all these lie in I. This shows that a.a.s. $\chi(Z_p, k) \leq 20g \leq O(k/\log p)$, completing the proof. $\square$

For $k = \Theta(p)$ the above estimate is tight up to a constant factor, as shown by the following, which is essentially proved in [14].

Proposition 3.4 *For any constant $1/2 \geq c_1 > 0$ there are two constants $b_1, b_2 > 0$ so that for $k = \lfloor c_1 p \rfloor$, a.a.s., $b_1 \frac{p}{\log p} \leq \chi(Z_p, k) \leq b_2 \frac{p}{\log p}$.*

The upper bound is proved in Theorem 2.1. The lower bound follows from the result of Green [14] who showed that the maximum size of an independent set in the relevant Cayley graph is, a.a.s., $O(\log p)$. Note that Green's proof actually deals with Cayley sum-graphs, rather than Cayley graphs, but the proof for Cayley graphs is analogous. Note also that he only deals with the case $k = p/2$, but his argument carries over to all admissible values of c_1. We omit the details.

3.2 General cyclic groups

The probabilistic proof in the previous section can be extended to general cyclic groups, by a more careful computation of the variance. As before, the main idea is showing that for a logarithmic number of random elements x_i of Z_n, with high probability there is a multiplier a so that ax_i lies in the middle third of Z_n for all i. This is done by proving the analog of Lemma 3.1 for Z_n, but the computation of the variance here requires some work. Luckily this has already been done in [7], and hence we will simply quote and apply a result from that paper.

Lemma 3.5 ([7], Corollary 4.1) *For every fixed $\alpha > 0$ and all $r > r_0(\alpha)$ the following holds. Let A be a set of r elements in the one dimensional torus $T = R/Z$, let N be a large integer, and let $(xA + y) (mod\ 1)$ be a random set of T, where x is a uniform random integer in $\{1, 2, \ldots, N\}$ and y is a random real in T. Let I be a fixed interval of length β in T, and let the random variable Y give the cardinality of $(xA + y) \cap I$. Then the expectation of Y is βr, and its variance, for all sufficiently large N, is at most $r^{1+\alpha}\beta^{1-\alpha}$.*

We apply the lemma where the set A consists of elements of the form $\frac{a}{n}$ for integers a. These elements represent the members of the cyclic group Z_n. As in the proof of Lemma 3.1 we can, given $1 > \delta > 2\mu > 0$, define a family of $\lceil \frac{1}{\mu} \rceil$ intervals in T, each of length $\delta - \mu$, so that any interval (with arbitrary real endpoints) in T of length δ fully contains one of these intervals. We can then repeat the proof of the lemma, replacing the expression for the variance by the estimate given in Lemma 3.5, noting that in our case, if N is any multiple of n, then the choice of a random uniform integer in $\{1, 2, \ldots, N\}$ is equivalent to a uniform choice in $\{1, 2, \ldots, n\}$, as we only multiply this number by fractions of the form a/n modulo 1. Interpreting the obtained elements as members of Z_n, this gives the following.

Lemma 3.6 *For any positive $\alpha > 0$, $\delta > 2\mu > 0$ and any $n > r > r_0(\alpha)$ the following holds. Let I be a cyclic interval in Z_n of size $|I| = \delta n$. Let $A \subset Z_n$ be an arbitrary subset of $r > r_0$ elements of Z_n, and let x be a uniformly chosen random element of Z_n. Define $A' = \{a \in A \ : \ xa \in I\}$ (where the product is computed in Z_n). Then the probability that the size of A' is smaller than $(\delta - 2\mu)|A|$ satisfies*

$$Pr(|A'| < (\delta - 2\mu)|A|) \leq \lceil \frac{1}{\mu} \rceil \frac{(\delta - \mu)^{1-\alpha} r^{1+\alpha}}{\mu^2 r^2} < \frac{2}{\mu^3 r^{1-\alpha}}.$$

The above lemma suffices to prove the following analogs of Theorem 3.2 and Theorem 3.3 with no essential change in the proofs, besides the obvious minor modifications required in the computation. We omit the details.

Theorem 3.7 *For any fixed $\epsilon > 0$, if n is an integer and $1 \leq k \leq (1 - \epsilon) \log_3 n$ then the chromatic number $\chi(Z_n, k)$ is, a.a.s., at most 3.*

Theorem 3.8 *For all k, $1 < k \leq n/2$, $\chi(Z_n, k) \leq O(1 + \frac{k}{\log n})$ a.a.s.*

The analog of Proposition 3.4 also holds for any cyclic group Z_n, where the lower bound is (essentially) proved in [14], and the upper bound follows from Theorem 2.1.

4 Abelian groups

4.1 Elementary abelian 2-groups

When $B = Z_2^t$ is an elementary abelian 2-group of order $n = 2^t$, and $k \leq 2.999t$ we can determine the typical chromatic number of (B, k) accurately. This is described in the following theorem, in which $\omega(1)$ denotes any positive function that grows to infinity, arbitrarily slowly, with t .

Theorem 4.1 *For $B = Z_2^t$ the following hold.*

(i) If $k \leq t - \omega(1)$ then $\chi(B, k) = 2$ a.a.s.

(ii) If $t + \omega(1) \leq k \leq 3t - \omega(1)\sqrt{t}$ then $\chi(B, k) = 4$ a.a.s.

(iii) If $k = t + \Theta(1)$ then the probability that $\chi(B, k) = 2$ as well as the probability that $\chi(B, k) = 4$ are both bounded away from 0.

Note that the chromatic number of (B, k) above is never 3. Indeed, a somewhat surprising known result of Payan asserts that no Cayley graph of Z_2^t can have chromatic number 3.

Lemma 4.2 (Payan [20]) *If the chromatic number of a Cayley graph of an elementary abelian 2-group is at least 3, then it is at least 4.*

Although the result sounds surprising, its elegant proof is not very difficult. It proceeds by showing that if such a graph contains an odd cycle, then it contains the graph of an even dimensional discrete cube together with additional edges connecting every pair of antipodal vertices, and by proving that the chromatic number of each of these graphs is 4, as they contain a so called generalized Mycielski graph. See [20] for more details.

It is not difficult to use Edmonds's well known result about covering matroids by bases (see [12]) in order to show that a random set of $2t$ elements of Z_2^t can be partitioned , a.a.s., into two linear bases. This can be used to show that for $k = 2t$ the chromatic number of (B, k) is, a.a.s., at most 4. In order to deal with somewhat higher values of k we prove the following.

Lemma 4.3 *Let S be a random set of at most $3t - \omega(1)\sqrt{t}$ elements of Z_2^t. Then, a.a.s., S can be partitioned into two disjoint sets S_1, S_2 so that S_2 is linearly independent and S_1 contains no subset of odd cardinality whose sum is the 0 vector.*

Proof. Let $S = (s_1, s_2, \ldots, s_k)$, where $k = 3t - g\sqrt{t}$, and g grows arbitrarily slowly to infinity as t grows to infinity, be a random sequence of elements of Z_2^t. Starting with both S_1, S_2 empty, examine these elements one by one. Whenever an s_i is not a sum of an even number of previous vectors placed already in S_1, add it to S_1. Else, put it in S_2. Note that by definition, the set S_1 produced at the end of this process will not contain any subset of odd cardinality with sum 0. We proceed to show that a.a.s. the process ends with S_2 of cardinality smaller than $t - (g/5)\sqrt{t}$ which is linearly independent.

During the process of producing S_1 and S_2 let us also produce a linear basis of S_1 consisting of all elements that when they are added to S_1 increase its rank. Thus, the basis consists of all members of S_1 which are not linear combinations of previous members of S_1. When exposing a new element of S, let us first examine only whether or not it is a linear combination of the current basis of S_1. If so, it is clearly a uniform linear combination, and let us next examine if it is a sum of an odd number of basis elements, or an even number (each of these events happens with probability $1/2$). If it is a sum of an even number-it is thrown into S_2-note that at this point it is a uniform sum of an even number of these basis elements, and we will expose the actual sum only once we consider S_2. If it is a sum of an odd number of the basis elements, it is added to S_1.

For any positive i, the probability that s_i is a linear combination of the previous elements $s_1, \ldots, s_{i-1}$ is at most 2^{i-1-t}, and hence a.a.s. all the first, say, $t - g$ members of S are placed in S_1. After that, at most g elements will be added to S_1 while increasing its rank. However, at least $k - t$ times we will have an element that is a linear combination of the previous basis elements of S_1, and whenever such a random element is exposed, it is a sum of an odd number of such elements with probability exactly $1/2$. Therefore, the number of extra elements added to S_1 in that part of the process is a binomial random variable with parameters $k - t$ and $1/2$ and hence a.a.s. its value is at least, say, $(k-t)/2 - (g/4)\sqrt{t}$, leaving at most $t - (g/4)\sqrt{t} + g < t - (g/5)\sqrt{t}$ elements for S_2. Note that S_1 contains a basis and additional elements each of which is a sum of an odd number of basis elements, and it is therefore clear that no sum of an odd number of members of S can be the zero vector (as altogether such a sum is a sum of an odd number of basis elements).

What about S_2 ? It consists of at most $t-(g/5)\sqrt{t}$ elements, and each of them is a sum of an even number of the basis elements of S_1. We also know that each such element is a uniform linear combination of an even number of members from some prefix consisting of at least $t-g$ basis elements of S_1. This means that each member of S_2 is uniformly distributed over a set of at least 2^{t-g-1} vectors. However, that implies that the probability that some given fixed subset of elements of S_2 has sum 0 is at most $2^{-(t-g-1)}$, since we can expose all members of this subset but the last one, and then the last one still has at least 2^{t-g-1} possibilities and at most one of them can make the sum 0. As there are only at most $2^{t-(g/5)\sqrt{t}}$ subsets of S_2 this implies that the probability that one of them adds to 0 is at most $2^{t-(g/5)\sqrt{t}} \cdot 2^{-(t-g-1)}$, which is negligible. This shows that a.a.s. S_2 is linearly independent, completing the proof. □

Proof of Theorem 4.1:

(i) Put $k=t-h$ with $h=\omega(1)$. The probability that all k vectors in S are linearly independent is at least $\prod_{i=0}^{k-1}(1-\frac{1}{2^{t-i}})$, as each term $(1-\frac{1}{2^{t-i}})$ is exactly the conditional probability that vector number $i+1$ does not lie in the span of the previous ones assuming all previous ones are linearly independent. This product is bigger than $1-\frac{1}{2^{t-k}}=1-\frac{1}{2^h}$, showing that a.a.s all vectors in S are linearly independent. Therefore, a.a.s. no nontrivial linear combination of members of S is 0, and in particular no sum of an odd number of members of S is 0. Since an odd cycle in the Cayley graph (Z_2^t, S) is exactly a sum of an odd number of members of S that add to 0 this shows that a.a.s. (B,k) contains no odd cycle and is thus bipartite. This proves (i).

(ii) By the obvious monotonicity it suffices to show that for $k=t+2h$, with $h=\omega(1)$, $\chi(Z_2^t,k)\geq 4$ a.a.s and that $\chi(Z_2^t, 3t-\omega(1)\sqrt{t})\leq 4$ a.a.s.

To prove that a.a.s. $\chi(Z_2^t, t+2h)\geq 4$ observe, first, that a.a.s, the first $t+h$ members of S span Z_2^t, as the probability they do not is at most $2^t\cdot 2^{-(t+h)}=2^{-h}$ since if they do not span Z_2^t all should be orthogonal to some vector in it. Assuming this is the case, fix a basis among the first $t+h$ members of S, and expose the last h vectors in S. Each of them is the sum of an even number of the basis elements we fixed with probability exactly 1/2. If it is, then this generates an odd cycle in the graph. The probability this fails to happen is at most 2^{-h}. This shows that a.a.s. the graph contains an odd cycle, and hence by Lemma 4.2 its chromatic number is at least 4.

To show that a.a.s. $\chi(Z_2^t, 3t-\omega(1)\sqrt{t})$ is at most 4 let S be a random sequence of elements of Z_2^t, $|S|=3t-\omega(1)\sqrt{t}$. Apply Lemma 4.3 to partition S into two sets S_1, S_2 as in the lemma. Such a partition exists a.a.s., and it provides a partition of the set of edges of the Cayley graph into two bipartite graphs, implying the assertion of (ii).

(iii) For $k=t-h$, $h=O(1)$ nonnegative, the probability that the first $t-h-1$ vectors in S are linearly independent is bounded away from zero, and then the conditional probability that the last vector is a sum of an even number of them, creating an odd cycle, is at least 2^{-h-2} which is bounded away from zero. By Lemma 4.2 if this happens then the chromatic number is at least 4.

Similarly for $k = t + h$, $h = O(1)$ positive, the probability that the first t vectors in S form a basis is bounded away from zero, and then the conditional probability that each of the last h vectors is a sum of an odd number of the basis elements is at least 2^{-h}, which is bounded away from zero. If this happens, then there is no odd cycle and the graph is 2-colorable. Monotonicity thus completes the proof. $\square$

For bigger values of k the situation is less clear. If $k \geq 2^{\Omega(t)}$ then by Theorem 2.1, part (i) we know that $\chi(Z_2^t, k) \leq O(\frac{k}{\log k})$ a.a.s., and by Theorem 2.1, parts (ii) and (ii) we get some lower bounds. Note that for $k = 2^{ct}$ with $0 < c < 1$ the gap between the upper and lower bounds is large. For smaller values of k, say, $2.99t < k < 2^{o(t)}$, we can show that $\chi(Z_2^t, k) \leq O(k/t)$ as follows. Put $p = \lceil \log_2(k/t) \rceil + 2$, and consider only the vectors in S whose first p coordinates are all 0. A.a.s. their number is smaller than, say, $t/2$, which is much smaller than their length (as $p = o(t)$). Thus a.a.s. they are linearly independent and the Cayley graph in which the only edges correspond to these elements has chromatic number 2. The Cayley graph corresponding to all other edges (arising from the members of S with nonzero values in the first p coordinates) can be trivially colored by 2^p colors- simply color each vertex by the vector of its first p-coordinates. The product coloring gives a proper coloring of our graph with at most $2^p \cdot 2 = O(k/t)$ colors.

We have thus proved the following simple proposition.

Proposition 4.4 *For all $t < k \leq 2^{t-1}$, $\chi(B, k) \leq O(k/t)$ a.a.s.*

For very large values of k one can get a sharper estimate, using the results of Green [14]. Indeed, these results give the following.

Proposition 4.5 *For every c, $0 < c \leq 1/2$ there are $b_1 = b_1(c), b_2 = b_2(c) > 0$ so that for $n = 2^t$ and $k = cn = c2^t$, a.a.s.*

$$b_1 \frac{n}{\log n \log \log n} \leq \chi(B, k) \leq b_2 \frac{n}{\log n \log \log n}.$$

Indeed Green proves in [14] that for $c = 1/2$ the largest independent set in $(Z_2^t, c2^t)$ is, a.a.s, of size $\Theta(\log n \log \log n)$, providing the lower bound for this case. Moreover, his proof shows that there is such an independent set consisting of all nonzero elements of a linear subspace, and we can thus color by the cosets of this subspace. His proof works essentially as it is for all other values of c which are bounded away from 0. Note that he considers Cayley sum graphs, but for Z_2^t the definitions of Cayley sum-graphs and Cayley graphs coincide.

4.2 General abelian Groups

For general abelian groups, it is not difficult to see that if k is small with respect to n, then the chromatic number is typically at most 3.

Theorem 4.6 *For any abelian group B of size n and any $k \leq \frac{1}{4} \log \log n$, the chromatic number $\chi(B, k)$ satisfies $\chi(B, k) \leq 3$ a.a.s.*

Proof. Let $B = Z_{n_1} \oplus Z_{n_2} \oplus \ldots \oplus Z_{n_r}$ be a general abelian group of order $n = n_1 n_2 \ldots n_r$. Let $S = \{s_1, .., s_k\}$ be a random subset, where $s_i = (s_{i1}, s_{i2}, \ldots, s_{ir})$ for $1 \leq i \leq k$. Note that each s_{ij} is a random uniform element of Z_{n_j}.

There is a natural graph-homomorphism from the Cayley graph (B, k) to the Cayley graph (Z_{n_j}, k), mapping each vertex to its j-th coordinate. Thus, the chromatic number of (B, k) is at most that of (Z_{n_j}, k) for every j (where here we define the chromatic number to be infinite if some generator vanishes in Z_{n_j}.)

If for some j $|Z_{n_j}| \geq \log n$, then the result follows from Theorem 3.7. Else, $r \geq \frac{\log n}{\log \log n}$. For a fixed j, the probability that all values s_{ij} for $1 \leq i \leq k$ fall into the open middle third of Z_{n_i} ensuring chromatic number at most 3, is at least, say, $(\frac{1}{4})^k \geq \frac{1}{\sqrt{\log n}}$. Thus, the probability that this fails for all values of j is at most $(1 - \frac{1}{\sqrt{\log n}})^r < e^{-\sqrt{\log n}/\log \log n}$. It follows that in this case, a.a.s. , the chromatic number of at least one graph (Z_{n_j}, k) is at most 3, and hence so is the chromatic number of (B, k). $\square$

5 Open problems

The general problem of determining or estimating more accurately the chromatic number of a random Cayley graph in a given group with a prescribed number of randomly chosen generators deserves more attention. It may be interesting, in particular, to study the case of the symmetric group S_n.

Regarding other groups, it seems plausible to believe that for every solvable group B of size n and every $k \leq 0.01 \log n$, $\chi(B, k) \leq 3$ a.a.s., but we have not been able to prove or disprove this statement.

Is it true that for every group B of size n and every $k \leq n/2$, the typical chromatic number $\chi(B, k)$ differs from the degree of regularity of (B, k) only by a poly-logarithmic factor (in n), that is: is $\chi(B, k) = \tilde{\Theta}_n(k)$ a.a.s. ?

Another interesting question is the study of the concentration of the chromatic number $\chi(B, k)$, that is, the standard deviation of this quantity. Our results show that for several families of groups this is $o(1)$ for small values of k (though for elementary abelian 2-groups Z_2^t and for $k = t + \Theta(1)$ the deviation is $\Theta(1)$, by Theorem 4.1). Since our model here does not fix the size of the set of generators (as we are choosing them with repetitions), there is a rather simple argument showing that in this model, for cyclic groups Z_n the standard deviation is a least $\Omega(\frac{\sqrt{n}}{\log n})$ for some values of $k = \Theta(n)$. If, however, we fix the degree of regularity $|S \cup S^{-1}|$ of the graph, the standard deviation may well be smaller.

Finally, it seems interesting to investigate systematically other invariants and properties of random Cayley graphs. For a finite group B and an integer k, the random Cayley graph (B, k) is a natural model of a random regular graph, and the study of its properties is often challenging. In this paper we focused on the investigation of its chromatic number, while some of the earlier papers mentioned here deal with its expansion and spectral properties. The problem of developing a general theory of this class of random graphs deserves further attention.

Acknowledgment I would like to thank Jarek Grytczuk for helpful comments.

References

[1] N. Alon, Large sets in finite fields are sumsets, J. Number Theory 126 (2007), 110–118.

[2] N. Alon, I. Kriz and J. Nešetřil, How to color shift hypergraphs, Studia Scientiarum Mathematicarum Hungarica 30 (1995), 1-11.

[3] N. Alon, M. Krivelevich and B. Sudakov, List coloring of random and pseudo-random graphs, Combinatorica 19 (1999), 453-472.

[4] N. Alon, M. Krivelevich and B. Sudakov, Coloring graphs with sparse neighborhoods, J. Combinatorial Theory, Ser. B 77 (1999), 73-82.

[5] N. Alon and V. D. Milman, λ_1, isoperimetric inequalities for graphs and superconcentrators, J. Combinatorial Theory, Ser. B 38(1985), 73-88.

[6] N. Alon and A. Orlitsky, Repeated communication and Ramsey graphs, IEEE Transactions on Information Theory 41 (1995), 1276–1289.

[7] N. Alon and Y. Peres, Uniform dilations, Geometric and Functional Analysis 2 (1992), 1–28.

[8] N. Alon and Y. Roichman, Random Cayley graphs and expanders, Random Structures and Algorithms 5 (1994), 271–284.

[9] J. Bourgain and A. Gamburd, Uniform expansion bounds for Cayley graphs of $SL_2(F_p)$, Ann. of Math. (2) 167 (2008), no. 2, 625–642.

[10] S. Czerwiński, Random runners are very lonely, to appear.

[11] Christofides, Demetres; Markstrom, Klas, Expansion properties of random Cayley graphs and vertex transitive graphs via matrix martingales. Random Structures Algorithms 32 (2008), no. 1, 88–100.

[12] J. Edmonds, Minimum partition of a matroid into independent subsets. J. Res. Nat. Bur. Standards Sect. B 69B (1965) 67–72.

[13] B. J. Green, Essay submitted for the Smith's Prize, Cambridge University, 2001.

[14] B. J. Green, Counting sets with small sumset, and the clique number of random Cayley graphs, Combinatorica 25 (2005), 307–326.

[15] A. J. Hoffman, On eigenvalues and colorings of graphs, B. Harris Ed., Graph Theory and its Applications, Academic, New York and London, 1970, 79-91

[16] Z. Landau and A. Russell, Random Cayley graphs are expanders: a simple proof of the Alon-Roichman Theorem, Electron. J. Combin. 11 (2004), Research Paper 62, 6 pp.

[17] P. S. Loh, and L. J. Schulman, Improved expansion of random Cayley graphs. Discrete Math. Theor. Comput. Sci. 6 (2004), no. 2, 523–528.

[18] A. Lubotzky, R. Phillips and P. Sarnak, Ramanujan graphs, Combinatorica 8(3) (1988), 261-277.

[19] G. A. Margulis, Explicit group-theoretical constructions of combinatorial schemes and their application to the design of expanders and superconcentrators, Problems of Information Transmission 24(1988), 39-46.

[20] C. Payan, On the Chromatic Number of Cube-like Graphs. Discrete Math., 103 (1992), 271–277.

[21] X. Zhu, Circular chromatic number: a survey, Discrete Mathematics, 229 (1-3) (2001), 371-410.