EXPANDER GRAPHS AND SIEVING IN COMBINATORIAL STRUCTURES

FLORENT JOUVE$^{\boxtimes}$ and JEAN-SÉBASTIEN SERENI

Communicated by B. Martin

Abstract

We prove a general large-sieve statement in the context of random walks on subgraphs of a given graph. This can be seen as a generalization of previously known results where one performs a random walk on a group enjoying a strong spectral gap property. In such a context the point is to exhibit a strong uniform expansion property for a suitable family of Cayley graphs on quotients. In our combinatorial approach, this is replaced by a result of Alon–Roichman about expanding properties of random Cayley graphs. Applying the general setting we show, for instance, that with high probability (in a strong explicit sense) random coloured subsets of integers contain monochromatic (nonempty) subsets summing to 0, and that a random colouring of the edges of a complete graph contains a monochromatic triangle.

1. Introduction

The relevance of using families of expander graphs for studying objects or solving problems coming from a broad variety of mathematical areas has been emphasized in numerous ways in recent years. Notably the combination of sieving arguments together with expansion properties has proved particularly efficient. Let us mention the ground-breaking work [4], where the mix of such techniques enabled the authors to detect almost primes in a variety of non-Abelian situations. A different kind of sieve together with the same expansion properties has also been exploited in the context of group theory [12] or to obtain quantitative results in the probabilistic Galois theory of arithmetic groups [8]. In the sieving processes used in the aforementioned works, one is naturally led to a crucial step where some *spectral gap* property is needed. Deep results about the groups involved then come into play. Typically the required properties are provided by recent breakthroughs in algebraic combinatorics that have led to strong forms of Lubotzky's property (τ) (so-called *superstrong approximation*, a culminating point of the study of which is the work by Golsefidy and Varjú [7]).

Indeed, the spectral gap result needed can be tautologically interpreted as a property of expansion of a certain family of graphs.

The goal of the present paper is to establish a general large-sieve inequality in a purely combinatorial setting. More precisely, we develop an axiomatic version of a sieve in the context of countable families of Cayley graphs, which are randomly generated via a random walk. This can be viewed as a generalization of the framework used for instance in [9, Ch. 7] or [8] to a situation where spectral gap properties on groups are no longer available. To circumvent the lack of algebraic structure, we will exploit the fact that random graphs are 'good expanders'. Precisely we will use a result of Alon–Roichman [3] according to which a family of random Cayley graphs is very likely to form a good family of expanders. However, further difficulties arise: contrary to the usual situation, expansion properties are not sufficient to ensure good enough cancellation in the correlation sums appearing. To obtain the required cancellation, we introduce structural properties and use concentration arguments. To the best of our knowledge, this point of view is new within sieving contexts and it seems to us that it might lead to new uses and theoretical studies of sieves in a combinatorial setting.

After describing our general setting and proving our main result (Theorem 2.3), we briefly present some concrete uses of our result to specific questions.

We describe several applications of our method to the study of typical properties of subgraphs of a given graph. To produce random elements for which we want to test if some given property holds, we perform a random walk on the family of graphs studied (cf. also [8]). Another approach could consist in quantifying the proportion of elements that satisfy an expected property among a finite subset of the family of graphs considered. For the applications we have in mind, this question would in fact be much easier. As a matter of fact, we do need to quantify proportions of 'good' elements as part of our sieving process.

The paper is organized as follows. In Section 2 we state and prove the main result and we emphasize the way in which Alon–Roichman's theorem enables us to work in a setting which is combinatorial in nature (whereas earlier works such as [9, Ch. 7] require a more algebraic framework). The rest of the paper is devoted to applications of the main result. We conclude the paper with remarks on further questions that may be of interest and that can be successfully investigated via our method. We notably state a Ramsey-type result (together with a sketch of the proof) obtained by suitably adapting the arguments used in the second application.

NOTATION 1.1. If X is a finite set, then $\#X$ and $|X|$ synonymously denote the cardinality of X.

If X is a finite graph, then $\mathrm{Adj}(X)$ is the adjacency operator sending a $\mathbf{C}$-valued function on the vertices of X to the function $(x \mapsto \sum_y f(y))$, where the sum is over the neighbours y of the vertex x. If X is moreover d-regular (that is, every vertex of X has degree d), then the *normalized adjacency operator* is $(1/d) \cdot \mathrm{Adj}(X)$.

If G is a group and $S \subset G$, then $X(G, S)$ is the Cayley graph on G with edge set

$$S \cup S^{-1} := \{s \in G : s \in S \text{ or } s^{-1} \in S\}.$$

If G is finite and Abelian, then $\hat{G}$ is the character group of G. If x is a nonnegative real number, then $\lceil x \rceil$ and $\lfloor x \rfloor$ are the least integer greater than or equal to x and the greatest integer smaller than or equal to x, respectively. If R is a positive integer, then $[R]$ is the set $\{1, \ldots, R\}$. Given a probability space $(\Omega, \Sigma, \boldsymbol{P})$ and two events A and B such that $\boldsymbol{P}(B) \neq 0$, we let $\boldsymbol{P}(A \mid B)$ be the *conditional probability* $\boldsymbol{P}(A \cap B)/\boldsymbol{P}(B)$. If f and g are two real-valued functions defined on a set $\mathcal{D}$ and depending on a set $\mathcal{P}$ of parameters, then $f(x) \ll_{\mathcal{P}_0} g(x)$ means that there exists a positive constant C depending only on the subset $\mathcal{P}_0 \subseteq \mathcal{P}$ such that $|f(x)| \leqslant C|g(x)|$ whenever $x \in \mathcal{D}$.

2. The general setting and a large sieve for graphs

2.1. Random walk large sieve: statement of the main result. Stating our main result requires some definitions and a precise description of the general setting. Let G be an Abelian group (in this section, the group law is noted multiplicatively) and $\Lambda \subset \mathbf{N}$ be a (nonnecessarily finite) set of indices. We suppose that we are given a family $(H_\ell)_{\ell \in \Lambda}$ of subgroups of G such that for each ℓ the index $n_\ell := [G : H_\ell]$ is finite. We let $\rho_\ell \colon G \to G/H_\ell$ be the canonical projection. If ℓ and ℓ' are two distinct elements of Λ, we define $\rho_{\ell,\ell'} \colon G \to G/H_\ell \times G/H_{\ell'}$ by $g \mapsto (\rho_\ell(g), \rho_{\ell'}(g))$.

We fix once and for all a probability space $(\Omega, \Sigma, \boldsymbol{P})$ and an arbitrarily small real $\delta \in (0, 1)$. One of the main ingredients needed in our analysis is a theorem of Christofides and Markström [5, Theorem 5], which is stated later (see Theorem 2.7). Keeping the notation of [5, Theorem 5], we define two quantities:

$$\psi(\delta) := 2((2 - \delta)\log(2 - \delta) + \delta \log \delta)^{-1} \tag{2.1}$$

and, for each $\ell \in \Lambda$,

$$\kappa(b_\ell, \ell; \delta) := \lceil \psi(\delta)(\log n_\ell + b_\ell + \log 2) \rceil,$$

where $b := (b_\ell)$ is a parameter (a sequence of positive real numbers).

Now let $s_1^{(\ell)}, \ldots, s_{\kappa(b_\ell,\ell;\delta)}^{(\ell)}$ be independent identically distributed random variables taking values in G/H_ℓ. The random walk on G we want to consider is obtained by lifting the sets $\{s_1^{(\ell)}, \ldots, s_{\kappa(b_\ell,\ell;\delta)}^{(\ell)}\}$ (and their 'inverses' so that all the graphs considered are then undirected) to G. To that purpose, we define the random variable

$$S_\ell(b_\ell, \delta) := \{s_1^{(\ell)}, \ldots, s_{\kappa(b_\ell,\ell;\delta)}^{(\ell)}\} \cup \{(s_1^{(\ell)})^{-1}, \ldots, (s_{\kappa(b_\ell,\ell;\delta)}^{(\ell)})^{-1}\},$$

which takes values in the set of subsets of G/H_ℓ. For every index $\ell \in \Lambda$ and every integer $m \in \{1, \ldots, \kappa(b_\ell, \ell; \delta)\}$, we need to choose a representative $\tilde{s}_m^{(\ell)} \in G$ of $s_m^{(\ell)}$. The particular representative we choose is imposed by the following condition of *admissibility*. (As we shall establish later on, for each family of subgroups there exists at most one admissible local sequence in the sense of Definition 2.1.)

DEFINITION 2.1. Let $(H_\ell)_{\ell \in \Lambda}$ be a fixed family of subgroups of finite index of G and, for each $\ell \in \Lambda$, let R_ℓ be a set of representatives of G/H_ℓ. The sequence $(H_\ell, R_\ell)_{\ell \in \Lambda}$ is called an *admissible local sequence* for G if:

(i) $\bigcap_{\ell \in \Lambda} H_\ell = \{1\}$; and
(ii) for all $\ell, \ell' \in \Lambda$, $\ell \neq \ell' \Rightarrow R_\ell \subseteq H_{\ell'}$.

Let us assume that the sequence (H_ℓ) of subgroups of G is such that there is an admissible local sequence (H_ℓ, R_ℓ) for G. For each ℓ and each $m \in \{1, \ldots, \kappa(b_\ell, \ell; \delta)\}$, we choose the unique representative $\tilde{s}_m^{(\ell)}$ of $s_m^{(\ell)}$ in R_ℓ. The aforementioned uniqueness of an admissible local sequence ensures (see Lemma 2.4) that this defines in a unique way elements $\tilde{s}_m^{(\ell)}$ in G. Next we set

$$\tilde{S}_\ell(b, \delta) := \{\tilde{s}_1^{(\ell)}, \ldots, \tilde{s}_{\kappa(b_\ell, \ell; \delta)}^{(\ell)}\} \cup \{(\tilde{s}_1^{(\ell)})^{-1}, \ldots, (\tilde{s}_{\kappa(b_\ell, \ell; \delta)}^{(\ell)})^{-1}\}.$$

The subset of G we use to perform a random walk on G is

$$S(b, \delta) := \prod_{\ell \in \Lambda}^{*} (\{1\} \cup \tilde{S}_\ell(b_\ell, \delta)). \tag{2.2}$$

Let us explain precisely what the notation means. If $A_1, \ldots, A_k$ are k subsets of G, the product $\prod_{i=1}^{k} A_i$ is the subset $\{a_1 \ldots a_k : a_i \in A_i\}$ of G. Here the symbol $\prod^*$ means that for all but finitely many of the ℓ the ℓth factor picked equals 1. Finally, $S(b, \delta)$ is not seen as a random variable but as the product over ℓ of elements either equal to 1 or picked in $\tilde{S}_\ell(b_\ell, \delta)$ evaluated at a common $\omega \in \Omega$. In other words, we fix once and for all an element ω of Ω; picking an element of $S(b, \delta)$ amounts to picking 1 or an element of some $\tilde{S}_\ell(b_\ell, \delta)(\omega)$ and then computing the product of these elements.

With notation as above, we perform the following random walk on G. It is defined in the same way as in [9, Ch. 7].

$$\begin{cases} X_0 = g_0, \\ X_{k+1} = X_k \xi_{k+1} \quad \text{for } k \geqslant 0, \end{cases}$$

where g_0 is a fixed element in G and the steps ξ_k are independent, identically distributed random variables with distribution

$$\boldsymbol{P}(\xi_k = s) = \boldsymbol{P}(\xi_k = s^{-1}) = p_s = p_{s^{-1}}$$

for every k and every $s \in S(b, \delta)$, and where $(p_s)_s$ is a sequence of positive real numbers indexed by $S(b, \delta)$ such that

$$\sum_{s \in S(b, \delta)} p_s = 1.$$

Of course the random walk depends on the parameters $b = (b_\ell)_\ell$ and δ. If Λ is finite, the most natural such random walk is certainly the one defined by uniformly distributing the steps, that is, $p_s := \#S(b, \delta)^{-1}$ for every $s \in S(b, \delta)$. In general, we require that the sum of probabilities p_s over elements $s \in S(b, \delta)$ that are mapped by $\rho_{\ell, \ell'}$ to any given $(s', t') \in S_\ell(b_\ell, \delta) \times S_{\ell'}(b_{\ell'}, \delta)$ is not too small. Precisely, we assume throughout the paper that for any given $L \geqslant 1$, every $\ell, \ell' \in \Lambda \cap [1, L]$ and every $(s', t') \in S_\ell(b_\ell, \delta) \times S_{\ell'}(b_{\ell'}, \delta)$,

$$\sum_{\substack{s \in S(b, \delta) \\ \rho_{\ell, \ell'}(s) = (s', t')}} p_s \geqslant \frac{1}{\kappa(b_L, L, \delta)^2}, \tag{$\star$}$$

which seems an intuitive generalization of the uniform distribution to a general set Λ.

REMARK 2.2. In the illustrative examples discussed in Section 3, the growth of $\kappa(b_\ell, \ell, \delta)$ is linear as a function of ℓ. Therefore, in our examples the lower bound in ($\star$) becomes $\gg L^{-2}$. Thus, we see that some choices of sequence (p_s) are not suitable for ($\star$) to hold. For instance, fix an arbitrarily small $\varepsilon > 0$ and define (p_s) in such a way that for $s = 1$ (the identity element of G) one has $p_1 = 1 - \varepsilon$ and such that $\sum_{s \in S(b,\delta) \setminus \{1\}} p_s = \varepsilon$. For such a sequence (p_s), we can certainly produce infinitely many pairs of distinct indices $\ell, \ell' \in \Lambda$ contradicting ($\star$) (for example for the choice $L = \max(\ell, \ell')$). The condition ($\star$) should therefore be seen as a guarantee that the distribution of the steps ξ_j (that is, the sequence (p_s)) is not 'too chaotic' or in other words that the series of general terms (p_s) does not converge too fast.

By studying the properties of the random walk $(X_k)_k$, our aim is to describe the behaviour of a 'generic' element $g \in G$. To do so, we make use of Kowalski's abstract large-sieve procedure extensively described, together with applications, in his book [9]. As in every sieve method, one can only handle cases where the typical properties at issue can be detected locally. To be more precise, we fix for each $\ell \in \Lambda$ a conjugacy invariant subset $\Theta_\ell \subset G/H_\ell$. The probability we want to upper bound is

$$\boldsymbol{P}(\forall \ell \in \Lambda, \, \rho_\ell(X_k) \notin \Theta_\ell).$$

When applicable, the method shall produce effective upper bounds for the probability with which X_k satisfies a fixed property that can be detected by the condition $\rho_\ell(X_k) \notin \Theta_\ell$ for some $\Theta_\ell \subset G/H_\ell$. Our main result is the following abstract sieve statement. We refer the reader to the book by Kowalski [9, Proposition 3.5] for a (self-contained) sieve statement that Theorem 2.3 builds on. For more information on the random walk sieve used here, see also [9, Ch. 7].

THEOREM 2.3. *With notation as above, we set $\kappa_N := \kappa(b_N, N, \delta)$ for every $N \in \mathbf{N}$ and*

$$C_0 := \sup_{L \in \mathbf{N}_{>0}} \max_{\substack{\ell \neq \ell' \in \Lambda \\ L \leqslant \ell, \ell' \leqslant 2L}} \frac{\# S_\ell(b_\ell, \delta)}{\# S_{\ell'}(b_{\ell'}, \delta)}.$$

Assume that condition ($\star$) holds. Then there exists a positive real ν such that for every positive integer k,

$$\boldsymbol{P}(\rho_\ell(X_k) \notin \Theta_\ell, \, \forall \ell \in \Lambda_L) \leqslant \boldsymbol{P}(C_0 = \infty) + \sum_{\ell \in \Lambda_L} e^{-b_\ell}$$

$$+ \left(1 + \left(\sum_{\ell \in \Lambda_L} n_\ell \right) (1 - \kappa_{2L}^{-2} \nu)^k \right) \left(\sum_{\ell \in \Lambda_L} \frac{\# \Theta_\ell}{n_\ell} \right)^{-1},$$

where L is any fixed positive integer $\Lambda_L := \Lambda \cap [L, 2L]$ and the constant ν depends only on C_0, the set $S(b, \delta)$ and the distribution of the steps ξ_j (that is, the sequence (p_s)).

In applications, the probability that C_0 is infinite will be very small via a suitable choice of parameters. Later on we prove a lemma (Lemma 2.9) the purpose of which is to bound efficiently $\boldsymbol{P}(C_0 = \infty)$. We end this section by proving the aforementioned uniqueness of admissible sequences, when they exist.

LEMMA 2.4. *Let $(H_\ell)_{\ell \in \Lambda}$ be a family of subgroups of finite index of G. There exists at most one family $(R_\ell)_{\ell \in \Lambda}$ where R_ℓ is a set of representatives of G/H_ℓ such that (H_ℓ, R_ℓ) is an admissible local sequence.*

PROOF. Let $(H_\ell, R_\ell^{(1)})$ and $(H_\ell, R_\ell^{(2)})$ be two admissible local sequences for G. Fix $\ell_0 \in \Lambda$ and let $r_1 \in R_{\ell_0}^{(1)}$ and $r_2 \in R_{\ell_0}^{(2)}$ be representatives of the same element of G/H_{ℓ_0}, so there exists $h \in H_{\ell_0}$ such that $r_1 = r_2 h$. For any $\ell \neq \ell_0$, condition (ii) of Definition 2.1 implies that $r_1 \in H_\ell$ and $r_2 \in H_\ell$ and therefore $h \in H_\ell$. Thus, $h \in H_{\ell'}$ for every $\ell' \in \Lambda$. Now condition (i) of Definition 2.1 implies that $h = 1$ and hence $r_1 = r_2$. This shows that $R_{\ell_0}^{(1)} = R_{\ell_0}^{(2)}$, thereby concluding the proof.   □

2.2. Cayley graphs on quotients and expansion. Let G be an Abelian group. We are interested in the properties of the Cayley graphs on the groups $(G/H_\ell)_{\ell \in \Lambda}$ with edges corresponding to the values taken by the random variables $s_i^{(\ell)}$ for $i \in \{1, \ldots, \kappa(b_\ell, \ell; \delta)\}$. These graphs are regular: the regularity equals the number of distinct values taken by the random variables s_i.

REMARK 2.5. The assumption that G is Abelian is unnecessary for most of the results of this section. It is used in a crucial way only in Lemma 2.8. However, since our main result and the applications we develop only involve Abelian groups we chose to stick to this case, making the exposition simpler.

Throughout the paper, if $\mathcal{G}$ is a k-regular graph, then the *eigenvalues of $\mathcal{G}$* are the eigenvalues of the normalized adjacency operator $k^{-1} \operatorname{Adj}(\mathcal{G})$. An eigenvalue λ is *nontrivial* if $|\lambda| \neq 1$. The *spectral gap* $\varepsilon(\mathcal{G})$ of $\mathcal{G}$ is defined to be $\min\{1 - |\lambda| : \lambda$ is a nontrivial eigenvalue of $\mathcal{G}\}$ (recall that the eigenvalue -1 occurs if and only if $\mathcal{G}$ is bipartite). We adopt the following definition for an expander graph, which *slightly differs from the standard one*.

DEFINITION 2.6. Let γ be a real number satisfying $0 < \gamma \leqslant 1/2$. A k-regular graph $\mathcal{G}$ is a *γ-expander graph* if the spectral gap of $\mathcal{G}$ is at least γ.

Note in particular that according to the definition a k-regular graph with spectral gap greater than $1/2$ is a γ-expander graph for any $\gamma \in (0, 1/2)$.

The reason for introducing the above setup is a theorem of Alon and Roichman [3, Theorem 1], which has been subsequently improved by Landau and Russell [10, Theorem 2] and Loh and Schulman [11, Theorem 1]. The last improvement obtained so far, which is the version we state and use, is due to Christofides and Markström [5, Theorem 5].

THEOREM 2.7 (Christofides–Markström). *With notation as above, fix an index ℓ in Λ. For every real number $\delta \in (0 , 1/2]$, the probability that $X(G/H_\ell, \{s_1^{(\ell)}, \ldots, s_{\kappa(b_\ell, \ell; \delta)}^{(\ell)}\})$ is not a δ-expander graph is less than e^{-b_ℓ}.*

The statement can be rephrased by saying that it is highly probable that the Cayley graph $X(G/H_\ell, \{s_1^{(\ell)}, \ldots, s_{\kappa(b_\ell, \ell; \delta)}^{(\ell)}\})$ is a δ-expander graph, the counterpart being that the edge set has very large cardinality. Note that the definition of an expander graph we use is not completely equivalent to the usual definition. However, it is a standard fact

that the (usual) expansion property and the spectral gap property are closely related notions (see, for instance, [6, Theorem 1.2.3]), which allows us to use our definition harmlessly for our purposes.

Kowalski [9, Ch. 7] successfully combined large-sieve techniques with expansion properties in the setting of random walks on arithmetic groups. We want to transpose this principle in a combinatorial setting. When adapting Kowalski's work a nontrivial issue comes from the fact that the expansion property crucial to us is not automatically stable under Cartesian product. Precisely, [9, Ch. 7] relies on the fact that if S is a symmetric generating system for $\mathrm{SL}_2(\mathbf{Z})$ and if $\pi_d \colon \mathrm{SL}_2(\mathbf{Z}) \to \mathrm{SL}_2(\mathbf{Z}/d\mathbf{Z})$ is the reduction modulo d map for some $d \geqslant 2$, then the whole family of Cayley graphs on $\mathrm{SL}_2(\mathbf{Z})/\ker \pi_d$ (with respect to the projection of S) indexed by *squarefree integers* is expanding. In fact, it would be enough to have the same result with an index set replaced by the set of positive integers that are squarefree and products of at most two primes. (However, for Kowalski's purposes, considering the primes as the index set would not be sufficient.) To obtain a suitable combinatorial analogue of this method, the forthcoming lemma is sufficient. It shows that expansion properties of Cayley graphs are preserved, albeit only imperfectly, when one takes the Cartesian product of two base groups. The expansion ratio guaranteed by the lemma is strong enough for our purposes. However, we refer the interested reader to [1] for a much more sophisticated method that does produce expander 'product Cayley graphs'. (Recall that the definition of 'expander graph' we use slightly differs from the standard one.)

LEMMA 2.8. *Let δ be a real number in $(0, 1/2]$. With notation as above, assume that $X(G, S)$ and $X(H, T)$ are δ-expander Cayley graphs on finite Abelian groups G and H (with edge sets defined by $S \subseteq G$ and $T \subseteq H$, respectively). Then, for every $(x_0, y_0) \in G \times H$ with $x_0^2 = 1 = y_0^2$, the Cayley graph $X(G \times H, (S \times \{y_0\}) \cup (\{x_0\} \times T))$ is a $((1 + \gamma)^{-1}\delta)$-expander graph, where*

$$\gamma := \max\left\{ \frac{|S \cup S^{-1}|}{|T \cup T^{-1}|}, \frac{|T \cup T^{-1}|}{|S \cup S^{-1}|} \right\}.$$

PROOF. For convenience, set $Y := (S \times \{y_0\}) \cup (\{x_0\} \times T)$, $S^* := S \cup S^{-1}$, $T^* := T \cup T^{-1}$ and $Y^* := Y \cup Y^{-1}$. The eigenfunctions of the normalized adjacency operator on $X(G \times H, Y)$ are of the form

$$(\chi, \tau) \colon (g, h) \mapsto \chi(g)\tau(h)$$

for characters $\chi \in \hat{G}$ and $\tau \in \hat{H}$. The corresponding eigenvalues are of the form

$$\lambda_{\chi,\tau} := \frac{1}{|S^*| + |T^*|} \sum_{(g,h) \in Y^*} \chi(g)\tau(h).$$

Since $x_0^2 = 1 = y_0^2$, we have $\chi(x_0) = \chi(x_0^{-1})$ and $\tau(y_0) = \tau(y_0^{-1})$, so the sum splits as follows:

$$(|S^*| + |T^*|)\lambda_{\chi,\tau} = \tau(y_0) \sum_{g \in S^*} \chi(g) + \chi(x_0) \sum_{h \in T^*} \tau(h). \tag{2.3}$$

FIGURE 1. The product graph $X(\mathbf{Z}_4 \times \mathbf{Z}_4, \{(1, \sigma), (\sigma, 1)\})$, where σ is a generating element of $\mathbf{Z}_4$.

We deduce that

$$|\lambda_{\chi,\tau}| \leqslant \frac{|S^*|}{|S^*| + |T^*|}\left|\frac{1}{|S^*|}\sum_{g \in S^*}\chi(g)\right| + \frac{|T^*|}{|S^*| + |T^*|}\left|\frac{1}{|T^*|}\sum_{h \in T^*}\tau(h)\right|.$$

If both χ and τ are nontrivial, then $|\lambda_{\chi,\tau}| \leqslant 1 - \delta$ since each of $X(G, S)$ and $X(H, T)$ is a δ-expander, where $\delta \in (0, 1/2]$. If χ is trivial and τ is nontrivial, we obtain instead

$$|\lambda_{\chi,\tau}| \leqslant 1 - \delta(1 + |S^*|/|T^*|)^{-1}$$

and hence the result by symmetry of the roles played by G and H.   □

To better comprehend Lemma 2.8, we give several examples, which also allow us to demonstrate the necessity of the hypotheses made and the optimality of the bound given. For a positive integer n, we let $\mathbf{Z}_n$ be the cyclic group of order n. Consider first the case where both G and H are $\mathbf{Z}_4$, with S and T each consisting of a generating element of $\mathbf{Z}_4$. Thus, the graphs $X(G, S)$ and $X(G, T)$ are isomorphic to the undirected cycle C_4 with four vertices. (Recall that $X(G, S) = X(G, S^*)$ by the definition.) The spectral gap of C_4 is 1. Now choose x_0 and y_0 to be the neutral elements of G and H, respectively. The hypothesis of Lemma 2.8 are thus satisfied. Note that $\gamma = 1$, so, according to this lemma, the graph $X := X(G \times H, (\{x_0\} \times T) \cup (S \times \{y_0\}))$ should have spectral gap at least $(1 + 1)^{-1} \cdot 1 = \frac{1}{2}$. To check this, observe that X is the 4-regular graph depicted in Figure 1: it consists of two disjoint cycles of size 8 the vertices of which are 'linked using cycles of length 4'. This graph indeed has spectral gap exactly $\frac{1}{2}$. This can actually be directly deduced from the proof of Lemma 2.8 by using (2.3), thereby obtaining a precise expression for the eigenvalues of the product graph. More generally, one deduces that performing the same construction as we just did but starting from $\mathbf{Z}_{2k}$ for any integer $k \geqslant 2$ yields an infinite family of examples where the bound given by Lemma 2.8 is attained, showing its optimality. (The spectral gap of the two (isomorphic) starting graphs will be $1 - |\cos(\pi(k + 1)/k)|$ and that of the product graph exactly half this quantity.)

The hypothesis that x_0 and y_0 must be elements with order at most 2 in their respective groups is necessary, as is seen by taking $G := \mathbf{Z}_3 \times \mathbf{Z}_5 = \langle \sigma \rangle \times \langle \tau \rangle$ and $H := \mathbf{Z}_6 = \langle \mu \rangle$ and setting $S := \{\sigma\}$, $T := \{\mu\}$, $x_0 := \sigma\tau$ and $y_0 := \mu^3$. (Thus, y_0 is of order 2 while x_0 is of order 15.) The graph $X(G, S)$ consists of five disjoint triangles while $X(H, T)$ is a cycle of length 6. Consequently, each of these graphs has spectral gap $\frac{1}{2}$. However, the spectral gap of the graph $X(G \times H, ((\{x_0\} \times T) \cup (S \times \{y_0\})))$ is less than 0.045, which is less than $\frac{1}{2} \cdot (1 + \gamma)^{-1} = \frac{1}{4}$.

In applications, it is important to keep control of the 'spectral gap loss', that is, the size of the parameter γ appearing in Lemma 2.8. To do so, we need the following technical lemma, which asserts in a precise quantitative way that it is harmless to suppose that the random variables $s_i^{(\ell)}$ take distinct values when evaluated simultaneously and as long as n_ℓ is fairly larger than $\kappa(b_\ell, \ell; \delta)$. This is a simple application of standard concentration principles.

LEMMA 2.9. *Keeping notation as above, fix $\ell \in \Lambda$. For readability, $\kappa(b_\ell, \ell; \delta)$ is abbreviated to κ_ℓ. Let X be the random variable that counts the number of distinct values in the multi-set $\{s_1^{(\ell)}, \ldots, s_{\kappa_\ell}^{(\ell)}\}$. One has:*

(1)   *if there exists a positive real ε, independent of ℓ, such that $1 - \kappa_\ell/n_\ell > \varepsilon$, then*

$$P(X < \kappa_\ell/2) \leqslant \exp(-\kappa_\ell \varepsilon^2/8);$$

(2)   *if $n_\ell \leqslant \kappa_\ell$, then*

$$P(X < n_\ell/2) \leqslant \binom{n_\ell}{\lceil n_\ell/2 \rceil} \cdot 2^{-\kappa_\ell}.$$

PROOF. Let $x_1, \ldots, x_{n_\ell}$ be the elements of G/H_ℓ. For each $i \in [n_\ell]$, let X_i be the 0–1 random variable that is equal to 1 if $x_i \in \{s_j^\ell : 1 \leqslant j \leqslant n_\ell\}$. Notice that $X = \sum_{i=1}^{n_\ell} X_i$. Consequently, the linearity of expectation implies that $E(X) = \sum_{i=1}^{n_\ell} E(X_i)$. Moreover, for each $i \in [n_\ell]$,

$$
\begin{aligned}
E(X_i) = 1 - P(X_i = 0) &= 1 - \left(1 - \frac{1}{n_\ell}\right)^{\kappa_\ell} \\
&\geqslant 1 - \exp(-\kappa_\ell/n_\ell) \\
&\geqslant \kappa_\ell/n_\ell - 1/2 \cdot (\kappa_\ell/n_\ell)^2,
\end{aligned}
$$

so that $E(X) \geqslant \kappa_\ell - 1/2 \cdot \kappa_\ell^2/n_\ell$.

Now, since X is determined by κ_ℓ independent trials and, for every possible outcome of the trials changing the outcome of any one trial can affect X by at most 1, the simple concentration bound [14, page 79] yields that for every positive number t,

$$P(|X - E(X)| > t) \leqslant 2 \exp\left(-\frac{t^2}{2\kappa_\ell}\right).$$

By symmetry of the random variable X around its mean, we know that for every positive number t,

$$P(E(X) - X > t) = \tfrac{1}{2} P(|E(X) - X| > t).$$

Therefore, to prove (1), one sets $t := \kappa_\ell/2 \cdot (1 - \kappa_\ell/n_\ell)$. This implies that

$$P(X < \kappa_\ell/2) \leq \exp\left(-\frac{\kappa_\ell}{8}(1 - \kappa_\ell/n_\ell)^2\right) \leq \exp\left(-\frac{\kappa_\ell \varepsilon^2}{8}\right).$$

To prove (2), one rather proceeds as follows. Notice that $X < n_\ell/2$ if and only if there exists a subset H' of G/H_ℓ of size $\lceil n_\ell/2 \rceil$ such that $\{s_i^\ell : 1 \leq i \leq \kappa_\ell\} \cap H' = \varnothing$. This happens with probability at most $2^{-\kappa_\ell}$. Consequently, we infer that

$$P(X < n_\ell) \leq \binom{n_\ell}{\lceil n_\ell/2 \rceil} \cdot 2^{-\kappa_\ell}. \qquad \square$$

2.3. Random walk large sieve: proof of the main result. We first state an easy consequence of the definition of $S(b, \delta)$ that is useful in our sieving procedure.

LEMMA 2.10. *For all distinct integers $\ell, \ell' \in \Lambda$,*

$$\rho_\ell(S(b; \delta)) = S_\ell(b_\ell, \delta) \cup \{1\}$$

and

$$\rho_{\ell,\ell'}(S(b, \delta)) = (S_\ell(b_\ell, \delta) \cup \{1\}) \times (S_{\ell'}(b_{\ell'}, \delta) \cup \{1\}),$$

where $\{1\}$ is the one-element subset containing only the trivial class in the relevant quotient group.

PROOF. We use condition (ii) in Definition 2.1: the image by ρ_ℓ of $S(b, \delta)$ is the product of elements all equal to 1 except maybe for the ℓ-factor, which can be any element of $\rho_\ell(\tilde{S}_\ell(b_\ell, \delta) \cup \{1\})$, that is, any element of $S_\ell(b_\ell, \delta) \cup \{1\}$.

The second equality is obtained using the same argument. $\qquad \square$

We now define one last piece of useful notation before starting the proof of Theorem 2.3. For indices ℓ and ℓ' in Λ, we set $G_{\ell,\ell'} := G/H_\ell \times G/H'_{\ell'}$ if $\ell \neq \ell'$ and $G_\ell(= G_{\ell,\ell'}) := G/H_\ell$ otherwise. The proof of Theorem 2.3 is based on an adaptation of that of [9, Proposition 7.2].

PROOF OF THEOREM 2.3. Fix a real number δ in $(0, 1/2]$ and let us split the probability we are interested in:

$$\begin{aligned} P(\forall \ell \in \Lambda_L, \rho_\ell(X_k) \notin \Theta_\ell) \leq\ & P(C_0 = \infty) \\ & + P(\exists \ell \in \Lambda_L, X(G/H_\ell, \rho_\ell(S(b, \delta))) \text{ is not a } \delta\text{-expander}) \\ & + P(E_0), \end{aligned} \tag{2.4}$$

where E_0 is the intersection of the three following events:

- $C_0 < \infty$;
- for each $\ell \in \Lambda_L$, the graph $X(G/H_\ell, \rho_\ell(S(b, \delta))$ is a δ-expander; and
- $\rho_\ell(X_k) \notin \Theta_\ell$.

As we shall see, the third summand can be bounded from above using sieving techniques. Moreover, the second summand can be handled by invoking Theorem 2.7. Indeed, since $\rho_\ell(S(b,\delta)) = S_\ell(b_\ell,\delta) \cup \{1\}$ by Lemma 2.10, we can show that the following statement holds (note that the statement would be trivial if we were only interested in edge expansion):

$$\boldsymbol{P}(\exists \ell \in \Lambda_L, \, X(G/H_\ell, \rho_\ell(S(b,\delta))) \text{ is not a } \delta\text{-expander})$$
$$\leqslant \boldsymbol{P}(\exists \ell \in \Lambda_L, \, X(G/H_\ell, S_\ell(b_\ell,\delta)) \text{ is not a } \delta\text{-expander}).$$

This inequality is a consequence of Lemma 2.11 that we state and prove at the end of this section. Applying Theorem 2.7 yields that

$$\boldsymbol{P}(\exists \ell \in \Lambda_L, \, X(G/H_\ell, S_\ell(b_\ell,\delta)) \text{ is not a } \delta\text{-expander}) \leqslant \sum_{\ell \in \Lambda_L} e^{-b_\ell}.$$

Let us now turn to the third summand of the right-hand side of (2.4). First, let B be the conditional probability that $\rho_\ell(X_k) \notin \Theta_\ell$ for each $\ell \in \Lambda_L$ conditioned on the facts that $C_0 < \infty$ and $X(G/H_\ell, \rho_\ell(S(b,\delta)))$ is a δ-expander for each $\ell \in \Lambda_L$. Notice that

$$\boldsymbol{P}(E_0) \leqslant B. \tag{2.5}$$

Now we are in a situation close to the axiomatic sieve method developed in [9]. We fix (nonnecessarily distinct) indices ℓ and ℓ' in Λ_L (defining $\rho_{\ell,\ell}$ to be ρ_ℓ) and a character λ of $G_{\ell,\ell'}$

$$\lambda \colon G \xrightarrow{\rho_{\ell,\ell'}} G_{\ell,\ell'} \xrightarrow{\lambda_0} \mathbf{C}^\times$$

factoring through $G_{\ell,\ell'}$ in such a way that λ_0 is a nontrivial character of $G_{\ell,\ell'}$.

We first prove the following statement. Assume that $C_0 < \infty$ and $X(G/H_\ell, \rho_\ell (S(b,\delta)))$ is a δ-expander for every $\ell \in \Lambda$. We assert that there exists a positive constant ν depending only on C_0, the set $S(b,\delta)$ and the distribution of the steps ξ_j (that is, the sequence (p_s)) such that

$$|\boldsymbol{E}(\lambda(X_k))| \leqslant (1 - \kappa_{2L}^{-2}\nu)^k. \tag{2.6}$$

Consider

$$M := \boldsymbol{E}(\lambda(\xi_k)) = \sum_{s \in S(b,\delta)} p(s)\lambda(s),$$

which is a well-defined element of $\mathbf{C}^\times$ since the series defining M converges absolutely. Let us also consider the complex number $M^+ := 1 - M$.

Note that M and M^+ are in fact real numbers since the set $S(b,\delta)$ as well as the distribution of the steps ξ_k are symmetric. We also need to define

$$N_0 := \boldsymbol{E}(\lambda(X_0)) = \sum_{t \in T} \boldsymbol{P}(X_0 = t)\lambda(t) \in \mathbf{C}^\times,$$

where T is a fixed (finite) subset of G containing the starting point g_0 of the random walk (X_k). (For simplicity, one can assume that $T = \{g_0\}$.)

The random variables X_0 and ξ_k being independent, it follows that for every positive integer k,

$$\boldsymbol{E}(\lambda(X_k)) = N_0 M^k.$$

We have $|N_0| \leqslant 1$ and we compute, first assuming that $\ell \neq \ell'$,

$$
\begin{aligned}
M^+ &= \sum_{s \in S(b,\delta)} p_s (1 - \lambda(s)) \\
&= \sum_{\substack{(s',t') \in S_\ell(b_\ell,\delta) \times S_{\ell'}(b_{\ell'},\delta)}} \left(\sum_{\substack{s \in S(b,\delta) \\ \rho_{\ell,\ell'}(s)=(s',t')}} p_s \right) (1 - \lambda_0(s',t')) \\
&\geqslant \frac{1}{\kappa(b_{2L}, 2L, \delta)^2} \min_{\psi \neq 1} \max_{(s',t') \in S_\ell(b_\ell,\delta) \times S_{\ell'}(b_{\ell'},\delta)} (1 - \psi(s',t')),
\end{aligned}
$$

where ψ runs over the nontrivial characters of $G_{\ell,\ell'}$. With the same notation, we deduce that for any $a \in S_\ell(b_\ell, \delta)$ and any $b \in S_{\ell'}(b_{\ell'}, \delta)$ both of order at most 2,

$$M^+ \geqslant \kappa(b_{2L}, 2L, \delta)^{-2} \min_{\psi \neq 1} \max_{(s',t') \in S_\ell(b_\ell,\delta) \times \{b\} \cup \{a\} \times S_{\ell'}(b_{\ell'},\delta))} (1 - \psi(s',t')).$$

Lemma 2.8 asserts that the family of Cayley graphs with vertex set $G_{\ell,\ell'}$ and edge set $S_\ell(b_\ell, \delta) \times \{b\} \cup \{a\} \times S_{\ell'}(b_{\ell'}, \delta))$ is a family of $(1 + C_0)^{-1}\delta$-expanders if the family $(X(G_\ell, S_\ell(b_\ell, \delta)))_{\ell \in \Lambda}$ is a family of δ-expanders. Thus, we can appeal to the translation of Lubotzky's property (τ) into the property of expansion of the corresponding Cayley graphs (see [13, Proposition 2.5]) to justify the existence of a positive constant $\nu(C_0, S(b,\delta), (p_s))$ which is uniform in $\ell, \ell' \in \Lambda$ and such that

$$M^+ \geqslant \kappa(b_{2L}, 2L, \delta)^{-2} \nu(C_0, S(b,\delta), (p_s)).$$

The case where $\ell = \ell'$ of that inequality is easier to establish and even holds under a stronger form in which the lower bound is independent of C_0. In fact, the argument goes in exactly the same way as in the case where $\ell \neq \ell'$, only this time we do not need to invoke Lemma 2.8. Also, we note that if one fixes an auxiliary index $\ell'' \in \Lambda_L \setminus \{\ell\}$, then condition $(\star)$ implies that

$$
\begin{aligned}
\sum_{\substack{s \in S(b,\delta) \\ \rho_\ell(s)=(s')}} p_s &= \sum_{t' \in S_{\ell''}(b_{\ell''},\delta)} \sum_{\substack{s \in S(b,\delta) \\ \rho_{\ell,\ell''}(s)=(s',t')}} p_s \\
&\geqslant \kappa(b_{2L}, 2L, \delta)^{-2} |S_{\ell''}(b_{\ell''}, \delta)| \\
&\geqslant \kappa(b_{2L}, 2L, \delta)^{-2}.
\end{aligned}
$$

To conclude the proof of the assertion, it suffices to observe that, because of our definition of expansion, the fact that the family $X(G_\ell, \rho_\ell(S(b,\delta)))$ is assumed to be a family of δ-expanders implies that these Cayley graphs are not bipartite and hence producing a lower bound for $1 + M$ is not required.

We can now finish the proof by using Kowalski's large-sieve inequality [9, Proposition 3.7]. Recalling the definition of B preceding (2.5),

$$B \leqslant \Delta(X_k; L)\left(\sum_{\ell \in \Lambda_L} \frac{\#\Theta_\ell}{n_\ell}\right)^{-1},$$

where one has the theoretical upper bound

$$\Delta(X_k; L) \leqslant \max_{\ell \in \Lambda_L} \max_{\chi \in \mathcal{B}_\ell^*} \sum_{\ell' \in \Lambda_L} \sum_{\chi' \in \mathcal{B}_{\ell'}^*} |W(\chi, \chi')|,$$

with

$$W(\chi, \chi') := \boldsymbol{E}(\chi \rho_\ell(X_k)\overline{\chi' \rho_{\ell'}(X_k)}) = \boldsymbol{E}([\chi, \overline{\chi'}]\rho_{\ell,\ell'}(X_k)).$$

Here for any $\ell \in \Lambda$ we let $\mathcal{B}_\ell$ be the character group of G/H_ℓ. We set further $\mathcal{B}_\ell^* := \mathcal{B}_\ell \setminus \{1\}$. Finally, if ψ_i is a character of a finite Abelian group G_i for $i \in \{1, 2\}$, then we let $[\psi_1, \psi_2]$ be the character $\psi_1 \otimes \psi_2$ of $G \times H$ if $G \neq H$ or of G otherwise.

In our setting, using [9, Lemma 3.4], we deduce that

$$[\chi, \overline{\chi'}]\rho_{\ell,\ell'} = \delta((\ell, \chi), (\ell', \chi'))\boldsymbol{1} + [\chi, \overline{\chi'}]_0 \rho_{\ell,\ell'},$$

where $\delta(\cdot, \cdot)$ is the Kronecker symbol and $[\chi, \overline{\chi'}]_0$ is the component of $[\chi, \overline{\chi'}]$ orthogonal to the trivial character $\boldsymbol{1}$. Thus, applying (2.6) to

$$\lambda := [\chi, \overline{\chi'}]_0 \rho_{\ell,\ell'},$$

$$|\boldsymbol{E}([\chi, \overline{\chi'}]_0 \rho_{\ell,\ell'}(X_k))| \leqslant (1 - \nu \kappa_{2L}^{-2})^k.$$

Putting everything together and using (2.5), we deduce as wished that

$$\boldsymbol{P}(E_0) \leqslant B \leqslant \left(1 + \left(\sum_{\ell \in \Lambda_L} n_\ell\right)(1 - \nu \kappa_{2L}^{-2})^k\right)\left(\sum_{\ell \in \Lambda_L} \frac{\#\Theta_\ell}{n_\ell}\right)^{-1}. \qquad \square$$

It remains to prove the following statement.

LEMMA 2.11. *Let G_0 be an Abelian group and let S be a subset of G_0. If $X(G_0, S)$ is a δ-expander graph, then so is $X(G_0, S \cup \{1\})$.*

PROOF. The statement is trivially true if $1 \in S$, so we assume that $1 \notin S$. Set $S^* := S \cup S^{-1}$ and $s^* := \#S^*$. Recall that Definition 2.6 implies that $\delta \in (0, 1/2]$. To prove the statement, it suffices to show that every nontrivial eigenvalue λ' of $X(G_0, S \cup \{1\})$ is such that $|\lambda'| \leqslant 1/2$ or $|\lambda'| \leqslant |\lambda|$ for some nontrivial eigenvalue λ of $X(G_0, S)$.

Let λ' be a nontrivial eigenvalue of $X(G_0, S \cup \{1\})$. Using the usual convention according to which a loop contributes 2 to the degree of a vertex, we deduce that $\lambda' = (2 + s^*)^{-1}(\sum_{s \in S^*} \chi(s) + \chi(1))$ for some nontrivial character χ of G_0. Therefore,

$$\lambda' = \frac{s^*}{2 + s^*}\lambda + \frac{1}{2 + s^*} = \lambda + \frac{1 - 2\lambda}{2 + s^*},$$

where $\lambda := (s^*)^{-1} \sum_{s \in S^*} \chi(s)$ is a nontrivial eigenvalue of $X(G_0, S)$.

Consequently, it is enough to prove that if $|\lambda'| > 1/2$, then $|\lambda'| \leqslant |\lambda|$. Suppose, on the contrary, that $|\lambda'| > 1/2$ and $|\lambda'| > |\lambda|$. Then $\lambda' > 0$. Indeed, otherwise $\lambda \leqslant -1/s^* < 0$ and hence $-\lambda + (2\lambda - 1/2 + s^*) > |\lambda| = -\lambda$ implies that $\lambda > 1/2$, which is a contradiction.

Hence, $\lambda' > 1/2$, which yields that $\lambda > 1/2$. However, this implies that $(1 - 2\lambda/2 + s^*) < 0$, so that $\lambda > \lambda' = |\lambda'|$, contrary to our assumption. This finishes the proof.   □

3. Illustrative examples

This section illustrates how Theorem 2.3 can be applied to various classical topics. Let us state two bounds on ψ that are useful in the forthcoming applications. An elementary study of the function ψ, which extends continuously to $[0\,,1/2]$, shows that ψ is increasing on that interval and therefore

$$1.442\ldots \simeq (1/\log 2) \leqslant \psi(\delta) \leqslant 4/\log(27/16) \simeq 7.644\ldots \qquad (3.1)$$

for any $\delta \in (0, 1/2]$. We now present some applications inspired by the classical Ramsey theory.

3.1. Towards a quantitative infinite Ramsey theory.
Our purpose is to illustrate how our method can be applied in the context of infinite Ramsey theory. Let us first recall the result we have in mind, established by Ramsey [15]. Given a set X and a nonnegative integer r, we define $X^{(r)}$ to be the collection of all subsets of X of cardinality r.

THEOREM 3.1 (Infinite Ramsey theorem [15]). *Let X be some countably infinite set. Let c and r be positive integers. Consider a given colouring $f : X^{(r)} \to \mathbf{Z}/c\mathbf{Z}$ of the elements of $X^{(r)}$ in c different colours. Then there exists some infinite subset A of X such that the function f is constant on $A^{(r)}$, that is, all subsets of A of cardinality r have the same image under f.*

For every function f, the *support* of f is the set of all elements e in the domain of f such that $f(e) \neq 0$. As in the statement of Ramsey's theorem, fix positive integers c and r. As our base set we choose $X := \mathbf{N}_{\geqslant 1}$. The set $\mathscr{C}^{(r)}$ of all possible c-colourings of subsets of cardinality r of X may be endowed with a group structure inherited from that of $\mathbf{Z}/c\mathbf{Z}$. Explicitly, the addition of two elements f and g of $\mathscr{C}^{(r)}$ is formally defined by

$$f + g : X^{(r)} \to \mathbf{Z}/c\mathbf{Z}, \quad A \mapsto f(A) + g(A).$$

The neutral element is the function that is identically 0.

We also fix an auxiliary positive integer j and set $\Lambda := \mathbf{N}_{\geqslant 1}$. We consider the subsets $I_\ell^{(r,j)} := \{(r + j)(\ell - 1) + 1, \ldots, (r + j)\ell\}$ of X indexed by $\ell \in \Lambda$. If j and r are fixed, then $I_\ell^{(r,j)}$ is an integral interval of size $r + j$ and different indices ℓ and ℓ' give rise to disjoint intervals $I_\ell^{(r,j)}$ and $I_{\ell'}^{(r,j)}$. For $\ell \in \Lambda$, let $E_\ell^{(r)}$ be the set of subsets of size r of $I_\ell^{(r,j)}$. Let C_ℓ be the collection of all colourings supported on $E_\ell^{(r)}$ and let H_ℓ be the

subgroup of all colourings of $\mathscr{C}^{(r)}$ supported on the complement of $E_\ell^{(r)}$ in $X^{(r)}$. In this way C_ℓ is a set of representatives for the quotient $\mathscr{C}^{(r)}/H_\ell$. Indeed, no two distinct functions in C_ℓ are congruent modulo an element of H_ℓ. Moreover, for any $f \in \mathscr{C}$, let f_C be the colouring equal to f on $E_\ell^{(r)}$ and equal to 0 everywhere else. It follows that $f_C \in C_\ell$ and $f - f_C \in H_\ell$ or equivalently $f \equiv f_C \pmod{H_\ell}$. Let $\rho_\ell \colon \mathscr{C}^{(r)} \to \mathscr{C}^{(r)}/H_\ell$ be the canonical surjection. The disjointness of the sets $I_\ell^{(r,j)}$ readily implies that (H_ℓ, C_ℓ) is an admissible local sequence for $\mathscr{C}^{(r)}$. In addition, we note that $|E_\ell^{(r)}| = \binom{r+j}{r}$. Summing up, we have thus established the following statement.

LEMMA 3.2. *The sequence $(H_\ell, C_\ell)_{\ell \in \Lambda}$ is an admissible local sequence for $\mathscr{C}^{(r)}$ and*

$$\forall \ell \in \Lambda, \quad n_\ell := (\mathscr{C}^{(r)} : H_\ell) = \#C_\ell = c^{\binom{r+j}{r}}.$$

We may now define on $\mathscr{C}^{(r)}$ a random walk (X_k) that satisfies the requirements of Theorem 2.3. We then ask the question:
at which speed do we reach a colouring X_k of the r-element subsets of X that exhibits a subset $A \subseteq X$ of size $r + j$, all the r-element subsets of which have the same colour?
The next statement answers that question.

THEOREM 3.3. *Let (X_k) be the random walk defined on $\mathscr{C}^{(r)}$ as in Section 2.3, with $\tilde{S}_\ell(b, \delta) \subseteq C_\ell$. Fix an integer $c \geqslant 2$, two positive integers j and r and a positive real number ε. For every positive integer k, let E_k be the event that no element of $\mathbf{N}^{(j+r)}$ has all its r-element subsets of the same colour in X_k. Then*

$$\forall k \geqslant 1, \quad \boldsymbol{P}(E_k) \ll k^{-1/2+\varepsilon},$$

where the implied constant depends only on ε, j, r, c, C_0, $S(b, \delta)$ and the sequence (p_s). This constant could be explicitly computed in terms of ε, j, r, c and the constant v of Theorem 2.3. As a function of j, this constant is unbounded.

PROOF. According to Lemma 3.2, we know that n_ℓ is independent of ℓ since $n_\ell = c^{\binom{r+j}{r}}$. Let us set $b_\ell := \ell$ for all $\lambda \in \Lambda$. Thus, $\kappa(b_\ell, \ell; \delta) \geqslant \psi(\delta)(\log n_\ell + \ell + \log 2)$, where ψ is defined by (2.1). Thus, $n_\ell \leqslant \kappa(b_\ell, \ell; \delta)$ if ℓ is large enough, for example if $\ell \geqslant L_0 := c^{\binom{r+j}{j}}$. Moreover, if for each $\ell \in \Lambda_L$ the set $S_\ell(b_\ell; \delta)$ contains at least $n_\ell/2$ distinct elements, then $C_0 \leqslant 2$. Therefore, applying Lemma 2.9(2), we deduce that for for every $L \geqslant L_0$,

$$\boldsymbol{P}(C_0 = \infty) \leqslant \boldsymbol{P}(\exists \ell \in \Lambda_L, \#S_\ell(b_\ell; \delta) < n_\ell/2)$$
$$\leqslant \phi_0(c, r, j) \cdot \sum_{\ell \in \Lambda_L} 2^{-\kappa(b_\ell, \ell; \delta)},$$

where one can take $\phi_0(c, r, j)$ to be $\binom{n_\ell}{\lceil n_\ell/2 \rceil}$, which depends only on c, r and j. Next, as $\ell \leqslant \kappa(b_\ell, \ell; \delta)$, we deduce that

$$\forall L \geqslant L_0, \quad \boldsymbol{P}(C_0 = \infty) \leqslant \phi_0(c, r, j) 2^{-L+1}.$$

Fix a positive integer k and a positive real number ε. For each $\ell \in \Lambda$, we set

$$\Theta_\ell := \{g \in \mathscr{C}^{(r)}/H_\ell : \text{the only representative of } g \text{ in } C_\ell \text{ is constant on } E_\ell^{(r)}\}.$$

Of course, $\#\Theta_\ell/n_\ell = c/n_\ell = c^{-\binom{r+j}{r}+1}$. Before going further, we note the existence of a constant $\psi_1(c, r, j)$ depending only on c, r and j such that $\kappa(n_\ell, b_\ell; \delta) \leqslant 8\ell + \psi_1(c, r, j)$. Lemma 3.2 ensures that the hypotheses of Theorem 2.3 are satisfied. Thus, abbreviating $\kappa_\ell(b_\ell, \ell; \delta)$ as κ_ℓ,

$$P(\rho_\ell(X_k) \notin \Theta_\ell, \; \forall \ell \in \Lambda_L) - P(C_0 = \infty)$$

$$\leqslant \left(1 + c^{\binom{r+j}{r}} |\Lambda_L| (1 - \nu \kappa_{2L}^{-2})^k\right) \left(\sum_{\ell=L}^{2L} \frac{|\Theta_\ell|}{n_\ell}\right)^{-1} + \sum_{\ell=L}^{2L} e^{-b_\ell}$$

$$\leqslant \left(1 + (L+1) \cdot c^{\binom{r+j}{r}} (1 - \nu \kappa_{2L}^{-2})^k\right) \frac{c^{\binom{r+j}{r}-1}}{L+1} + e^{1-L} - e^{-2L}$$

$$\leqslant \frac{c^{\binom{r+j}{r}}}{L} + c^{2\binom{r+j}{r}} (1 - \nu(16L + \psi_1(c, r, j))^{-2})^k,$$

where we use the inequalities $e^{1-x} - e^{-2x} \leqslant 1/x$ if $x > 0$ and $c^{\binom{r+j}{r}-1} + 1 \leqslant c^{\binom{r+j}{r}}$ that holds as soon as $c \geqslant 2$, $r \geqslant 0$, $j \geqslant 1$. (For the first inequality, define $f_1(x) = e - e^{-x}$ and $f_2(x) = e^x/x$ for $x > 0$. The function f_1 is increasing and one-to-one from $(0, \infty)$ to $(e - 1, e)$ whereas f_2 has a global minimum at $x = 1$, as is obviously seen by computing the derivative of f_2. Therefore, $f_2(x) \geqslant f_2(1) = e$ for all $x > 0$. We deduce that $f_1(x) \leqslant f_2(x)$ for $x > 0$. The inequality asserted follows immediately.) For any fixed $\varepsilon > 0$, set $L := \lceil k^{1/2-\varepsilon} \rceil$. For this to be compatible with the condition $L \geqslant L_0$, it is enough that $k^{1/2-\varepsilon} \geqslant c^{\binom{r+j}{r}}$. This inequality can be made to hold by modifying the implied constant in the estimate proven, thereby making it depend also on ε. We compare the order of magnitude of the two summands of the above right-hand side with the upper bound obtained for $P(C_0 = \infty)$. First, one has

$$P(C_0 = \infty) \ll_{L_0} k^{-1/2+\varepsilon}.$$

Also, since $L \geqslant 1$, we know that $16L + \psi_1(c, r, j) \leqslant (16 + \psi_1(c, r, j))L$, so

$$(1 - (16L + \psi_1(c, r, j))^{-2}\nu)^k = \exp\left(-\frac{\nu}{(16 + \psi_1(c, r, j))^2} k^{2\varepsilon} + O(\nu k^{-1+4\varepsilon})\right)$$

$$\ll \exp\left(-\frac{\nu}{(16 + \psi_1(c, r, j))^2} k^{2\varepsilon}\right),$$

with an absolute implied constant. We thus obtain the upper bound

$$\varphi(\varepsilon, r, j, c, C_0, S(b, \delta), (p_s))k^{-1/2+\varepsilon}$$

for the probability investigated, where $\varphi(\varepsilon, r, j, c, C_0, S(b, \delta), (p_s))$ is a positive constant depending only on the tuple $(\varepsilon, r, j, c, C_0, S(b, \delta), (p_s))$. This finishes the proof. □

REMARK 3.4. We point out an important limitation to our approach: we cannot dispense with the use of the auxiliary parameter j. More precisely, letting j tend to infinity in the inequality of Theorem 3.3 yields only a trivial upper bound for the probability investigated (this comes from the unboundedness of the implied constant in Theorem 3.3 as a function of j).

3.2. Monochromatic solutions to equations. It also seems relevant to study solutions of an equation through the perspective of Ramsey theory: can one destroy the solutions of an equation by partitioning the different values the variables can take? We are interested in the following question, which turns out to fit our setting.

Given a random c-colouring of a random subset A of $\mathbf{Z}$, what is the probability that A contains a monochromatic nonempty subset summing to 0?

We study this question in two steps. First, we leave aside colourings and just bound the probability that a random subset of $\mathbf{Z} \setminus \{0\}$ contains no subset summing to 0. To this end, the group G considered is that of all subsets of $\mathbf{Z} \setminus \{0\}$ with the symmetric difference Δ as group law. We then show how easily one can add constraints on colourings to this setting, by just considering the product of the group G with the group of all c-colourings of $\mathbf{Z} \setminus \{0\}$. So, in our setting, the coloured version essentially reduces to the first question.

Let G be the group consisting of all subsets of $\mathbf{Z} \setminus \{0\}$ endowed with the symmetric difference. For each positive integer ℓ (that is, we choose $\Lambda := \mathbf{N}_{\geqslant 1}$), we set $I_\ell := \{-\ell, \ell\}$, $C_\ell := 2^{I_\ell}$ and we define H_ℓ to be the subgroup of G consisting of all subsets of $\mathbf{Z} \setminus \{0\}$ that are disjoint from I_ℓ. Thus, C_ℓ forms a set of representatives for $G_\ell := G/H_\ell$. In particular, $n_\ell := [G : H_\ell] = 4$ and $(H_\ell, C_\ell)_{\ell \geqslant 1}$ is an admissible local sequence for G (since the sets I_ℓ are pairwise disjoint).

We set
$$\Theta_\ell := \left\{ X \in G_\ell : \forall \tilde{X} \in G, \rho_\ell(\tilde{X}) = X \Rightarrow \sum_{x \in \tilde{X}} x = 0 \right\},$$

so Θ_ℓ is a singleton, the unique element of which is represented by I_ℓ. Now one can define a random walk (X_k) as in Section 2.3. This random walk readily satisfies the requirements of Theorem 2.3. Further, observe that if a subset S of $\mathbf{Z}$ does not contain a nonempty subset summing to 0, then neither does the intersection of S with any fixed subset. Thus, the probability P_k that X_k does not contain a nonempty subset summing to 0 is at most
$$\mathbf{P}(\rho_\ell(X_k) \notin \Theta_\ell, \forall \ell).$$

We choose $b_\ell = \ell$ for all $\ell \in \Lambda$ and apply the same method as in the proof of Theorem 3.3 to bound $\sum_{\ell \in \Lambda_L} b_\ell$ from above. Since $|\Theta_\ell|/n_\ell = \frac{1}{4}$ for each positive integer ℓ, Theorem 2.3 implies that for every positive integer L and every positive integer k,

$$P_k - \mathbf{P}(C_0 = \infty) \leqslant \frac{1}{L} + (1 + (L+1) \max_{L \leqslant \ell \leqslant 2L} |G_\ell|(1 - \nu \kappa_{2L}^{-2})^k) \cdot \frac{4}{L}$$
$$= \frac{5}{L} + 8(1 - \nu \kappa_{2L}^{-2})^k.$$

We have $\kappa(b_\ell, \ell; \delta) \geqslant \ell$ and thus $\kappa(b_\ell, \ell; \delta) \geqslant n_\ell$ whenever $\ell \geqslant 4$. Applying Lemma 2.10, we obtain in the same way as before,

$$\mathbf{P}(C_0 = \infty) \leqslant 6 \sum_{\ell=L}^{2L} 2^{-\ell}$$

for any choice of $L \geqslant 4$. Observing that $\kappa_{2L} \leqslant 32L$ by (3.1), and setting $L := \lceil k^{1/2-\varepsilon} \rceil$, we compute as in the proof of Theorem 3.3,

$$(1 - \nu \kappa_{2L}^{-2})^k \ll_\nu \exp\left(-\frac{\nu k^{2\varepsilon}}{32^2}\right).$$

Consequently, we infer the following statement.

THEOREM 3.5. *Let (X_k) be a random walk on G defined as in Section 2.3 using $S(b,\delta)$, with $\tilde{S}_\ell(b,\delta) \subseteq 2^{I_\ell}$. Then for all $\varepsilon > 0$ there exists a positive constant C_ε (that can be computed explicitly as a function of ε and ν) depending only on ε, $S(b,\delta)$, C_0 and the sequence (p_s) such that for every positive integer k,*

$$P(X_k \text{ does not contain a nonempty subset summing to } 0) \leqslant C_\varepsilon k^{-1/2+\varepsilon}.$$

Let us now see how to deal with the coloured version, that is, we want to upper bound the probability that our random c-coloured subset does not contain a *monochromatic* nonempty subset summing to 0, where c is an integer greater than 1. It suffices to work in the product group $G := (2^{\mathbf{Z}\setminus\{0\}}, \Delta) \times \{f : \mathbf{Z} \setminus \{0\} \to \mathbf{Z}/c\mathbf{Z}\}$. For each positive integer ℓ (that is, we choose $\Lambda := \mathbf{N}_{\geqslant 1}$), the subgroup H_ℓ is defined to be

$$2^{\mathbf{Z}\setminus(I_\ell \cup \{0\})} \times \{f : \mathbf{Z} \setminus \{0\} \to \mathbf{Z}/c\mathbf{Z} : f(-\ell) = f(\ell) = 0\},$$

where I_ℓ is $\{-\ell, \ell\}$ as before.

Thus, $n_\ell := [G : H_\ell] = 4 \cdot 2^c = 2^{c+2}$, which does not depend on ℓ. A set of representatives for $G_\ell := G/H_\ell$ is $2^{I_\ell} \times \mathscr{F}_\ell$, where

$$\mathscr{F}_\ell := \{f : \mathbf{Z} \setminus \{0\} \to \mathbf{Z}/c\mathbf{Z} \ : \ f|(\mathbf{Z} \setminus (I_\ell \cup \{0\})) = 0\}.$$

Again, since the sets I_ℓ are pairwise disjoint, the sequence $(H_\ell, \mathscr{F}_\ell)$ is an admissible local sequence for G.

Defining Θ_ℓ to be $\{I_\ell\} \times \{f : \mathbf{Z} \setminus \{0\} \to \mathbf{Z}/c\mathbf{Z} \ : \ f \text{ is constant}\}$, it follows that $|\Theta_\ell|/n_\ell = c2^{-c-2}$. Since the hypotheses of Theorem 2.3 are satisfied, one obtains the following statement. (The proof goes along the same lines as that of Theorem 3.5—in particular, we choose $b_\ell = \ell$ and, for any fixed $\varepsilon > 0$, we set $L := \lceil k^{1/2-\varepsilon} \rceil$. Details are omitted.)

THEOREM 3.6. *Let (X_k) be a random walk on G defined as in Section 2.3 using $S(b,\delta)$, with $\tilde{S}_\ell(b,\delta) \subseteq 2^{I_\ell} \times \mathscr{F}_\ell$. For every positive integer k, let E_k be the event that X_k does not contain a monochromatic nonempty subset summing to 0. Then, for every positive real number ε and every positive integer k,*

$$P(E_k) \ll_\varepsilon k^{-1/2+\varepsilon},$$

where the implied constant could be explicitly computed as a function of ε, c and ν (with notation as in Theorem 2.3) and depends only on ε, c, $S(b,\delta)$, C_0 and the sequence (p_s).

At this point, it seems relevant to also discuss Ramsey theory for graphs.

3.3. Looking for monochromatic triangles. We let $\mathscr{G}$ be the (countable) infinite complete graph, that is, the graph with vertex set $\mathbf{N}$ in which every two distinct positive integers are neighbours. We fix an integer $c \geqslant 3$ and we define $\mathscr{C}$ to be the collection of all functions from the edges of $\mathscr{G}$ to $\mathbf{Z}/c\mathbf{Z}$. As earlier, the set $\mathscr{C}$ is naturally endowed with a group structure inherited from that of $\mathbf{Z}/c\mathbf{Z}$.

We are interested in monochromatic substructures of a given fixed size that may arise. To avoid unnecessary notation and abstraction, we shall focus on finding monochromatic triangles—though our strategy could be adapted effortlessly to the question of detecting monochromatic r-cliques or r-cycles for $r \geqslant 3$.

We define a family of subgroups $(H_\ell)_{\ell \in \Lambda}$ of $\mathscr{C}$, where $\Lambda := \mathbf{N}$. Consider a partition *in finite parts* $(I_\ell)_{\ell \in \Lambda}$ of Λ. We set $i(\ell) := |I_\ell|$ for $\ell \in \Lambda$. Let $E_\ell := \{(a, b) \in I_\ell^2 : a \neq b\}$, that is, E_ℓ is the set of all edges of $\mathscr{G}$ with both end vertices contained in I_ℓ. We define C_ℓ to be the collection of all functions $f \in \mathscr{C}$ with support contained in E_ℓ. Then H_ℓ is the collection of all functions $f \in \mathscr{C}$ such that $f|E_\ell \equiv 0$.

Let us give the necessary properties that the quotients $\mathscr{C}_\ell := \mathscr{C}/H_\ell$ satisfy.

LEMMA 3.7. *For each $\ell \in \Lambda$, the following holds:*

(i) C_ℓ *is a set of representatives of the quotient $\mathscr{C}_\ell$ and $(H_\ell, \mathscr{C}_\ell)$ is an admissible local sequence for $\mathscr{C}$; and*

(ii)   *the index of H_ℓ in $\mathscr{C}$ is $n_\ell := [\mathscr{C} : H_\ell] = |C_\ell| = c^{i(\ell)(i(\ell)-1)/2}$.*

PROOF.

(i) No two distinct functions in C_ℓ are congruent modulo an element of H_ℓ. Moreover, for any $f \in \mathscr{C}$, let f_C be the function equal to f on E_ℓ and equal to 0 everywhere else, that is, $f_C|E_\ell := f|E_\ell$ and $f_C|(E(\mathscr{G}) \setminus E_\ell) := 0$. It follows that $f_C \in C_\ell$ and $f - f_C \in H_\ell$ or, equivalently, $f \equiv f_C \pmod{H_\ell}$. Finally, the fact that (H_ℓ, C_ℓ) is an admissible local sequence for $\mathscr{C}$ is a consequence of the disjointness of the sets I_ℓ.

(ii) By the definition, $|E_\ell| = i(\ell)(i(\ell) - 1)/2$. The conclusion follows.   □

A practical way to rephrase part of the proof of Lemma 3.7 is to say that for each fixed integer ℓ in Λ and each element f of $\mathscr{C}$, the unique element in C_ℓ congruent to f modulo H_ℓ is the function equal to f on E_ℓ and to 0 outside of E_ℓ.

From now on, we assume that $i(\ell) \geqslant 3$ for $\ell \in \Lambda$. For each integer $\ell \in \Lambda$, let Θ_ℓ be the set of classes $\bar{f} \in \mathscr{C}_\ell$ such that the unique representative f of $\bar{f}$ in C_ℓ (the existence of which is asserted by Lemma 3.7) contains a monochromatic triangle in E_ℓ. In other words, $f \in \Theta_\ell$ if and only if I_ℓ contains three integers i_1, i_2 and i_3 such that $f((i_1, i_2)) = f((i_1, i_3)) = f((i_2, i_3))$. Observe that $|\Theta_\ell|/|\mathscr{C}_\ell| \geqslant c^{-2}$. Indeed, any function that restricts to a constant map (with values in $\mathbf{Z}/c\mathbf{Z}$) on a fixed triangle contained in E_ℓ surjects to an element of Θ_ℓ via ρ_ℓ.

Assume that δ is a fixed real number in $(0, 1/2]$. We set $b_\ell := \ell$. In particular, note that

$$\kappa(b_\ell, \ell; \delta) = \left\lceil \psi(\delta) \cdot \left(\frac{i(\ell)(i(\ell) - 1) \log c}{2} + \ell + \log 2 \right) \right\rceil.$$

Given $f^{(\ell)} \in S_\ell(b_\ell, \delta)$, we define $\tilde{f}^{(\ell)}$ to be its canonical representative in $\mathscr{C}$, that is, $\tilde{f}^{(\ell)} \in C_\ell$. In this context, the outcome of Theorem 2.3 is the following result.

PROPOSITION 3.8. *Let (X_k) be a random walk on $\mathscr{C}$ defined as in Section 2.3 using $S(b, \delta)$ (see (2.2)), with $\tilde{S}_\ell(b, \delta) \subseteq C_\ell$. For every positive integer k, let E_k be the event that X_k does not contain a monochromatic triangle. Then, with notation as in Theorem 2.3, one has for any fixed positive integers L and k,*

$$P(E_k) \leqslant P(C_0 = \infty) + \frac{c^2 + 1}{L} + 2c^{(1/2) \cdot i(2L)(i(2L)-1)+2}(1 - \kappa_{2L}^{-2}\nu)^k.$$

PROOF. Fix a positive integer k. Lemma 3.7 ensures that the hypotheses of Theorem 2.3 are satisfied. We obtain, applying Theorem 2.3,

$$P(\rho_\ell(X_k) \notin \Theta_\ell, \forall \ell \in \Lambda_L) - P(C_0 = \infty)$$

$$\leqslant \left(1 + \left(\sum_{\ell \in \Lambda_L} n_\ell\right)(1 - \kappa_{2L}^{-2}\nu)^k\right)\left(\sum_{\ell \in \Lambda_L} \frac{|\Theta_\ell|}{n_\ell}\right)^{-1} + \sum_{\ell=L}^{2L} e^{-b_\ell}$$

$$\leqslant e^{1-L} - e^{-2L} + (1 + (L+1) \cdot c^{i(2L)(i(2L)-1)/2}(1 - \kappa_{2L}^{-2}\nu)^k) \cdot \frac{c^2}{L}$$

$$\leqslant \frac{c^2 + 1}{L} + 2c^{(1/2) \cdot i(2L)(i(2L)-1)+2}(1 - \kappa_{2L}^{-2}\nu)^k,$$

where we used that $e^{1-x} - e^{-2x} \leqslant 1/x$ for $x \geqslant 1$. □

Different choices of sets I_ℓ may correspond to different speeds of rarefaction of nontypical structures. (We note, however, that the random walk itself does depend on the choice made for the sets I_ℓ.) More precisely, one can put additional constraints on the structure of the monochromatic triangles, for instance we may impose the three vertices to be consecutive integers as in the following theorem.

THEOREM 3.9. *With notation as in Proposition 3.8, for every positive integer k, let E'_k be the event that X_k does not contain a monochromatic triangle on three consecutive vertices. One has for every positive real number ε and for every positive integer k,*

$$P(E_k) \leqslant P(E'_k) \ll_\varepsilon k^{-1/2+\varepsilon},$$

where the implied constant can be computed explicitly (as a function of ε, c and ν (see Theorem 2.3)) and depends only on ε, c, $S(b, \delta)$, C_0 and the sequence (p_s).

PROOF. Set $I_\ell := \{3\ell - 2, 3\ell - 1, 3\ell\}$ for each $\ell \in \Lambda$. In particular, $i(\ell)(i(\ell) - 1) = 6$. Let us evaluate $P(C_0 = \infty)$. One has $\#S_\ell(b_\ell; \delta) \leqslant n_\ell = c^3$ and $\kappa(b_\ell, \ell; \delta) = \lceil \psi(\delta)(3 \log c + \ell + \log 2) \rceil \geqslant \ell$, by (3.1). In particular, $n_\ell \leqslant \kappa(b_\ell, \ell; \delta)$ for all $\ell \geqslant c^3$. Moreover, if we assume that for all $\ell \in \Lambda_L$ the set $S_\ell(b_\ell; \delta)$ contains at least $n_\ell/2 = c^3/2$ distinct elements, then $C_0 \leqslant 2$. Therefore,

$$P(C_0 = \infty) \leqslant P(\exists \ell \in \Lambda_L, \#S_\ell(b_\ell; \delta) < n_\ell/2)$$

$$\leqslant \binom{c^3}{\lceil \frac{c^3}{2} \rceil} \sum_{\ell \in \Lambda_L} 2^{-\kappa(b_\ell, \ell; \delta)}$$

for all $L \geqslant c^3$, by virtue of Lemma 2.9(2).

For $\varepsilon > 0$ fixed, set $L := \lceil k^{1/2-\varepsilon} \rceil$. For this to be compatible with the condition $L \geqslant c^3$, we need to have $k^{1-2\varepsilon} \geqslant c^6$. This inequality can be made to hold by modifying the implied constant in the estimate to be proven. As in the proof of Theorem 3.3,

$$\boldsymbol{P}(C_0 = \infty) \ll_c k^{-1/2+\varepsilon}.$$

Moreover, Proposition 3.8 implies that

$$\boldsymbol{P}(E'_k) \leqslant \boldsymbol{P}(C_0 = \infty) + \frac{c^2 + 1}{k^{1/2-\varepsilon}} + 2c^5(1 - \nu\kappa_{2L}^{-2})^k.$$

To find an upper bound for the third summand, we first use the assumption $L \geqslant c^3$ to deduce $\kappa_{2L} \leqslant 46L$ and then we compute

$$(1 - \nu\kappa_{2L}^{-2})^k = \exp\left(k\left(-\frac{\nu}{46^2 k^{1-2\varepsilon}}\right) + O(\nu k^{-2+4\varepsilon})\right) \ll_{c,\nu} \exp\left(-\frac{\nu}{46^2}k^{2\varepsilon}\right),$$

which finishes the proof.   □

We note that the contribution from the nonstandard case (that is, $X(G/H_\ell, S_\ell(b_\ell; \delta))$ is not an expander) is the probability with highest order of magnitude (among the three summands in the upper bound of Theorem 2.3) given our choice of parameters in the proof of the theorem. It is natural to compare Theorem 3.9 with what is known from Ramsey theory; this discussion is deferred to the next section.

4. Remarks and further applications

As mentioned earlier, the main purpose of our work is to obtain a general sieve statement in a purely combinatorial setting. Regarding the illustrative applications, the general line of thought is to give, for the intricate notion of randomness defined, explicit upper bounds for probabilities that we expect to be small.

Let us underline some peculiarities of the application developed in Sections 3.2 and 3.3. For monochromatic substructures, it follows from Ramsey's theorem [15] that for every fixed positive integer c, there exists an integer N such that if $n \geqslant N$, then every c-colouring of the edges of the complete graph K_n on n vertices contains a monochromatic triangle. Alon and Rödl [2] established that the smallest such integer N is $\Theta(3^c)$ as n tends to infinity (that is, there exist two constants ρ and ρ' such that for sufficiently large n, this value belongs to $[\rho \cdot 3^c, \rho' \cdot 3^c]$). In our setting, although the infinite complete graph is involved, only finite subgraphs of it are checked for the existence of monochromatic triangles. These subgraphs are not necessarily large enough for Ramsey's theorem to apply. In addition, we only consider monochromatic triangles with vertices contained in some prescribed set I_ℓ.

Another feature of the applications presented is uniformity of the decay rate with respect to the number c of colours involved. Actually, we even claim control of the dependency of the implied constant as a function of c, since this implied constant could be explicitly computed. No such uniformity holds in the context of Ramsey theory.

Indeed, as already mentioned, Alon and Rödl's theorem [2] asserts that the number of required vertices for Ramsey's theorem to hold grows exponentially fast with c.

Next let us comment on the common decay rate, roughly $1/\sqrt{k}$, in our various applications. When applying Theorem 2.3, we always have to find an upper bound of the rough form

$$\sum_{L \leqslant \ell \leqslant 2L} 2^{-\ell} + c_1 \sum_{L \leqslant \ell \leqslant 2L} 2^{-c_2 \ell} + \left(1 + c_3 L \left(1 - \frac{c_4 \nu}{L^2}\right)^k\right) \frac{c_5}{L},$$

where each parameter c_i is an absolute constant.

The fast decay of the first two summands is not an issue as soon as L is chosen to be roughly equal to some power of k. However, in the third summand one has to have simultaneously $L \to \infty$ and $(1 - c_4 \nu L^{-2})^k \to 0$ as $k \to \infty$. These constraints justify the choice $L = \lceil k^{1/2-\varepsilon} \rceil$ in all our applications. There is certainly room for improvement here (for example by choosing a different value for b_ℓ, rather than setting b_ℓ to be ℓ, or by modifying the sieve itself so that n_ℓ is not necessarily bounded as a function of ℓ), but we feel that ensuring the decay of the third summand will remain a rather serious constraint in general.

We highlight a strategy similar to that used in Section 3.2 that allows one to check for monochromatic arithmetic progressions for which the length, the common difference and the 'shape' are prescribed. Fix positive integers s (the desired length of the arithmetic progression), q (the desired common difference) and $c \geqslant 3$ (the number of colours). Similarly as before, let $\mathscr{C}$ be the group of all c-colourings of $\mathbf{N}$. We consider the subsets $I_\ell := \{\ell s q, \ell s q + q, \dots, \ell s q + (s-1)q\}$ for $\ell \in \Lambda := \mathbf{N}$. (It is this choice of particular subsets of $\mathbf{N}$ of length at least s that provides a control on the 'shape' of the arithmetic progressions to be found.) In this setting our method yields the following result.

THEOREM 4.1. *Let (X_k) be a random walk on $\mathscr{C}$ defined as in Section 2.3 using $S(b,\delta)$ via the admissible local sequence (H_ℓ, C_ℓ). For every positive real number ε and every positive integer k,*

$$\mathbf{P}(X_k \text{ contains no monochromatic arithmetic progression}$$

$$\text{with common difference } q \text{ and length } s) \ll k^{-1/2+\varepsilon},$$

where the implied constant could be computed explicitly as a function of c, s, q and ν (see Theorem 2.3 for the definition of ν) and depends only on (c, s, q), C_0, $S(b,\delta)$ and the sequence (p_s).

Let us sketch briefly the proof. For each $\ell \in \mathbf{N}$, let H_ℓ be the set of all functions $f : \mathbf{N} \to [c]$ such that $f|_{I_\ell} \equiv 0$. The index in $\mathscr{C}$ of each of these subgroups is c^s. Moreover, there is a collection of natural representatives C_ℓ for the classes modulo H_ℓ, namely the functions with support contained in I_ℓ. Thus, $n_\ell = c^{\#I_\ell}$ is independent of ℓ and, since the intervals C_ℓ are pairwise disjoint, the sequence (H_ℓ, C_ℓ) is an admissible local sequence for $\mathscr{C}$. Let Θ_ℓ be the set of classes modulo H_ℓ

whose unique representative in C_ℓ contains a monochromatic arithmetic progression of length s that is contained in I_ℓ. Then one has $|\Theta_\ell|/n_\ell \geqslant c^{-s}$.

Again we may apply Theorem 2.3 with $b_\ell = \ell$ for all $\ell \in \Lambda$. Similarly as before, $\boldsymbol{P}(C_0 = \infty)$ can be bounded from above: if $L \geqslant c^{(s-1)q+1}$, then $\kappa(b_\ell, \ell, \delta) \geqslant n_\ell$ whenever $\ell \geqslant L$ and hence Lemma 2.9(2) yields that $\boldsymbol{P}(C_0 = \infty) \leqslant \binom{c^s}{\lceil c^s/2 \rceil} \sum_{\ell \in \Lambda_L} 2^{-\ell}$. Therefore, $\boldsymbol{P}(C_0 = \infty) \leqslant \binom{c^s}{\lceil c^s/2 \rceil}/L$.

By Theorem 2.3, the probability that in X_k no monochromatic arithmetic progression with common difference q and length s is contained in I_ℓ, for all ℓ in Λ_L, is at most

$$\frac{\binom{c^s}{\lceil c^s/2 \rceil}}{L} + \frac{1}{L} + (1 + (L+1)c^{(s-1)q+1}(1 - \nu\kappa_{2L}^{-2})^k)(c^s L)^{-1}.$$

Since this last probability is, for every L, an upper bound on the probability that there is no monochromatic arithmetic progression in X_k with common difference q and length s, Theorem 4.1 follows by setting, for any fixed $\varepsilon > 0$, $L := \lceil k^{-1/2+\varepsilon} \rceil$.

We conclude by pointing out the following: van der Waerden's theorem [16] ensures that, for each fixed positive integer s and each integer $c \geqslant 3$, there exists an integer N such that if $n \geqslant N$ then any c-colouring of $[n]$ yields a monochromatic arithmetic progression of length s. In the above setting, we impose two additional conditions: the common difference of the arithmetic progression and a constraint on its form (it must be contained in one of the sets I_ℓ). Van der Waerden's theorem does not guarantee the existence of such an arithmetic progression and the aforementioned inequality is essentially an explicit lower bound on the speed of rarefaction of the colourings that do not yield a monochromatic arithmetic progression with the required properties. Furthermore, and as mentioned in the remarks about Sections 3.2 and 3.3, the uniformity of the decay rate with respect to the number of colours c is a quite interesting by-product of our approach.

References

[1] N. Alon and A. Lubotzky, 'Semi-direct product in groups and zig–zag product in graphs: connections and applications (extended abstract)', in: *42nd IEEE Sympos. Foundations of Computer Science (Las Vegas, NV, 2001)* (IEEE Computer Society, Los Alamitos, CA, 2001), 630–637.

[2] N. Alon and V. Rödl, 'Sharp bounds for some multicolor Ramsey numbers', *Combinatorica* **25**(2) (2005), 125–141.

[3] N. Alon and Y. Roichman, 'Random Cayley graphs and expanders', *Random Structures Algorithms* **5**(2) (1994), 271–284.

[4] J. Bourgain, A. Gamburd and P. Sarnak, 'Affine linear sieve, expanders, and sum-product', *Invent. Math.* **179**(3) (2010), 559–644.

[5] D. Christofides and K. Markström, 'Expansion properties of random Cayley graphs and vertex transitive graphs via matrix martingales', *Random Structures Algorithms* **32**(1) (2008), 88–100.

[6] G. Davidoff, P. Sarnak and A. Valette, *Elementary Number Theory, Group Theory, and Ramanujan Graphs*, London Mathematical Society Student Texts, 55 (Cambridge University Press, Cambridge, 2003).

[7] A. S. Golsefidy and P. P. Varjú, 'Expansion in perfect groups', *Geom. Funct. Anal.* **6**(22) (2012), 1832–1891.

[8]   F. Jouve, E. Kowalski and D. Zywina, 'Splitting fields of characteristic polynomials of random elements in arithmetic groups', *Israel J. Math.* **193**(1) (2013), 263–307.

[9]   E. Kowalski, *The Large Sieve and its Applications: Arithmetic Geometry, Random Walks and Discrete Groups*, Cambridge Tracts in Mathematics, 175 (Cambridge University Press, Cambridge, 2008).

[10]   Z. Landau and A. Russell, 'Random Cayley graphs are expanders: a simple proof of the Alon–Roichman theorem', *Electron. J. Combin.* **11**(1) (2004), Research Paper 62, 6.

[11]   P.-S. Loh and L. J. Schulman, 'Improved expansion of random Cayley graphs', *Discrete Math. Theor. Comput. Sci.* **2**(6) (2004), 523–528 (electronic).

[12]   A. Lubotzky and C. Meiri, 'Sieve methods in group theory I: powers in linear groups', *J. Amer. Math. Soc.* **25**(4) (2012), 1119–1148.

[13]   A. Lubotzky and A. Zuk, 'On property (τ)', Preprint, 2002.

[14]   M. Molloy and B. Reed, *Graph Colouring and the Probabilistic Method*, Algorithms and Combinatorics, 23 (Springer, Berlin, 2002).

[15]   F. P. Ramsey, 'On a problem of formal logic', *Proc. Lond. Math. Soc. (1)* **S2-30** (1930), 264–286.

[16]   B. L. van der Waerden, 'Beweis einer Baudetschen Vermutung', *Nieuw Arch. Wiskd.* **15**(1) (1927), 212–216.

FLORENT JOUVE, IMB, Université de Bordeaux, Talence, France
e-mail: florent.jouve@math.u-bordeaux.fr

JEAN-SÉBASTIEN SERENI, C.N.R.S. (LORIA), Vandœuvre-lès-Nancy, France
e-mail: sereni@kam.mff.cuni.cz