ON MIXING OF CERTAIN RANDOM WALKS, CUTOFF PHENOMENON AND SHARP THRESHOLD OF RANDOM MATROID PROCESSES

Igor Pak, Van H. Vu

Department of Mathematics
Yale University
New Haven, CT 06520

December 14, 1999

Abstract. In this paper we define and analyze convergence of the *geometric random walks*, which are certain random walks on vector spaces over finite fields. We show that the behavior of such walks is given by certain random matroid processes. In particular, the mixing time is given by the expected stopping time, and the cutoff is equivalent to sharp threshold.

We also discuss some random geometric random walks as well as some examples and symmetric cases.

Introduction

In the past two decades there has been an explosion in the number of applications of probabilistic method. In this paper we use probabilistic method to analyze a special case of Markov chains we call *geometric random walks*.

Here is a general setup of the problem. Let G be a finite group, and let S be a set of generators of G. Consider a Markov chain X_t on G which starts at the identity $X_0 = e$ and moves by the rule $X_{t+1} = X_t \cdot s$, where $s \in S$ is a random generator. It is easy to see that (under mild conditions) *"after a while"* the walk will be at an approximately uniform group element. The problem, however, is to quantify and compute how long is "after a while". This time is usually called *mixing time*. It depends in a complicated way on the the random walk and is normally very hard to estimate even in nice examples (see [AF,D1]). There is a large literature dedicated to finding bounds on mixing time as well to comparison of different definitions of mixing time (see [AF,LW]).

Suppose now we have a sequence of groups $\{G_i\}$ and their generating sets $\{S_i\}$, where $i \in \mathbb{N}$. One can try to quantify how rapidly the walks moves from the state of being "far from mixing" to the state of being "well mixed". Aldous and Diaconis observed (see [AD,D1]) that in many natural cases this transition happens in a period of time small compared to the mixing time. This is called *cutoff phenomenon*

Key words and phrases. Random walks, Markov chains, random graphs, stopping time, separation distance, cutoff phenomenon, sharp threshold.

and is somewhat analogous to the phase transition in statistical physics (see [D2]). While proven in many examples using asymptotically sharp estimates on the mixing time (see [D2]), the cutoff phenomenon remains a mystery yet to be solved.

In this paper we introduce a notion of a geometric random walk. Let $V = \mathbb{F}_q^n$ be an abelian group of vectors in a finite vector space. A subset $S \subset V$ is called *geometric* if with each $s \in S$ we have $a \cdot s \in S$ for all $a \in \mathbb{F}_q$. In other words, S must be a union of lines in V. Now let X_t be a random walk on V generated by S. We call it a *geometric random walk*. Examples of geometric random walks include random walks on a cube (the case when S forms a basis in $V = \mathbb{F}_2^n$) and other familiar combinatorial walks.

Consider now a vector matroid M corresponding to S. Define a *random matroid process* as follows. Start with an empty set and add random matroid elements one by one until we get a base. In this paper we show:

1) The mixing time of a geometric random walk is equal to the expected run time of the corresponding random matroid process (see sections 2, 4).

2) The cutoff for a geometric random walk exists if and only if the corresponding random matroid process has a sharp threshold (see section 8).

3) The cutoff exists if S is chosen randomly in a certain sense (see sections 7, 10, 11).

4) The cutoff can be proved in several cases (see sections 9, 12).

5) The expectation can be computed exactly in several natural cases (see sections 3, 5, 6).

Perhaps the main point of the paper is methodological. We establish connection between cutoff phenomenon for mixing of random walks, and sharp threshold for random matroids and graphs, (also known as phase transitions, as well as $0-1$ laws.)

While we do succeed in proving cutoff in several new cases, it is rather more important for us that we give a new look at the subject. Thus we go at length to reprove the cutoff in cases when it was long known (such as random walk on a cube). We hope this new approach can be useful in other cases, and perhaps help uncover the mystery of the phenomenon in a noncommutative situation which we do not touch in this paper.

Acknowledgements

We are grateful to Persi Diaconis for the introduction to the subject. We would also like to thank László Lovász, Gregory Margulis, and Richard Stanley for the interest in this work. Special thanks to Martin Hildebrand for some suggestions on style.

Part of the research was done when the first author was an NSF Postdoctoral Fellow at MIT.

1. Basic definitions

Let V be a d-dimensional space over the finite field $\mathbb{F}_q$, and let $O \in V$ be the origin. Denote $[k] = \{1, \dots, k\}$. Also, if $u_1, \dots, u_k$ are vectors in V, denote by $\langle u_1, \dots, u_k \rangle \subset V$ their linear span.

Let $A = \{v_1, \dots, v_m\} \subset V$ be a set of vectors in V such that $\langle v_1, \dots, v_m \rangle = V$. Define a *geometric random walk* $\mathcal{W}(A)$ to be a Markov chain X_t on vectors in V, such that $X_0 = O$ and

$$X_{t+1} = X_t + a(t) \cdot v_{i(t)}$$

where $a(t) \in \mathbb{F}_q$ and $i(t) \in [m]$ are uniform and independent random variables. One can think of X_t as a symmetric random walk on an abelian group $\mathbb{F}_q^n$ generated by elements $a \cdot v_i$, $i \in [m]$.

Consider an example. Suppose $q = 2$, $m = d$ and $A = \{v_1, \dots, v_m\}$. Then $\mathcal{W}(A)$ is equivalent to a lazy random walk on a cube $\mathbb{Z}_2^m$ which is defined by the following rule:

- *Choose a coordinate direction $i \in [m]$ uniformly. Flip a fair coin. If heads, move along that direction and if tails stay put.*

This walk was analyzed in a number of papers (see e.g. [D1,DGM,P1]). Roughly, the walks mixes after $O(n \log n)$ steps. The problem is in many ways similar to the coupon collector's problem (see [F,D1]). We will give a careful analysis of this walk in section 3 where the connection is made precise.

Denote by Q^k the probability distribution of the walk after k steps:

$$Q^k(v) = P(X_k = v)\,, \;\; v \in V$$

Observe that the Markov chain X_t is irreducible, aperiodic and reversible (see e.g. [F,AF]). Thus it is ergodic and the Q^k converges to a uniform stationary distribution $U \equiv 1/q^n$ as $k \to \infty$.

There are several ways to quantify how fast Q^k converges to U. The most commonly used are the *variation distance*

$$tv(k) = \max_{B \subset V} \big|Q^k(B) - U(B)\big| = \frac{1}{2} \sum_{v \in V} \left| Q^k(v) - \frac{1}{N} \right|$$

and the *separation distance*

$$s(k) = N \cdot \max_{v \in V} \left(\frac{1}{N} - Q^k(v) \right)$$

where $N = |V| = q^n$ is the total number of vectors in V.

For random walks on groups both distances have a similar asymptotic behavior, but the latter will suit better for our purposes. The separation distance has nice *submultiplicativity property*

$$s(m + k) \le s(m) \cdot s(k)\,, \text{where } m, k > 0.$$

Note also that $s(0) = 1$ and $tv(k) \le s(k)$ for all $k > 0$ (see [AD,AF,D1]).

Often it is useful to define a *mixing time* which is a single measure of the convergence. Again, there are several different measures which include but are not exhausted by the following two:

$$n_{1/2} = \min\big\{i : s(i) \le \frac{1}{2}\big\} = \min\big\{i : P^i(v) \ge \frac{1}{2N} \text{ for all } v \in V\big\}$$

and

$$\xi = 1 + s(1) + s(2) + \dots$$

The latter is called the *total separation* and the submultiplicativity property implies that $\xi < \infty$. It has same order of magnitude as $n_{1/2}$:

$$\xi \le n_{1/2} \le 2\xi$$

(see [P1]) and will be the main object of our study.

It is convenient to consider a generation function for the separation distances

$$\xi(z) = 1 + s(1) \cdot z + s(2) \cdot z^2 + \dots$$

which is called *separation series*. Clearly, $\xi = \xi(1)$. The function $\xi(z)$ is known to be rational in z and has no poles inside a unit disc $|z| \le 1$ (see [P1]).

It turns out that in case of the geometric random walks one can find an explicit combinatorial formula for the separation series and the total separation. This is done in the next section.

2. Strong uniform times

Let X_t be a finite Markov chain on a state space V with uniform stationary distribution. A *stopping rule* is an algorithm which observes the chain and stops it depending on the state passed. Denote by τ the *stopping time* of this rule. By $\varrho = X_\tau$ denote the *stopping state*. We think of τ and ϱ as of random variables.

The stopping time τ is called *strong uniform* if ϱ is uniform and independent of τ. In other words,

$$P(\varrho = v \,|\, \tau = k) = \frac{1}{|V|} \text{ for all } v \in V,\, k > 0$$

The main application of the strong uniform time is the following result.

Theorem 2.1 *Let τ be a strong uniform time for $\mathcal{M}$. Then for $k > 0$ we have:*

$$s(k) \le P(\tau > k), \quad \textit{and} \quad \xi \le E(\tau).$$

The first part is due to Aldous and Diaconis (see [AD,D1]) and the second part follows from the first part. In this form it is due to the first named author (see [P1]).

A strong uniform time τ is called *perfect* if $\xi = E(\tau)$. It is known that a perfect time always exists and must satisfy $s(k) = P(\tau > k)$ for all $k > 0$ (see [P1,D1]).

State $v \in V$ is called *halting* for a stopping time τ if the Markov chain always stops whenever it gets there. If a strong uniform time has a halting state, then it is perfect (see [P1,DF]). Therefore if we have a construction of a perfect time, we can immediately compute the separation distance $s(k)$ for all k.

Now we present a simple construction of a perfect time for a geometric random walk. Recall that $V \simeq \mathbb{F}_q^n$ and we have a fixed set of vectors $A = \{v_1, \ldots, v_m\} \subset V$. Our walk starts at the origin $X_0 = O$ and is defined as follows:

$$X_{t+1} = X_t + a \cdot v_i$$

where $a = a(t) \in \mathbb{F}_q$ and $i = i(t) \in [m]$ are uniform and independent random variables.

Consider a stopping rule which stops the walk as soon as vectors $v_{i(1)}, v_{i(2)}, \ldots$ generate the whole space V. We claim that this defines a strong uniform time. Formally, let $\tau = k$ be the first time we have $\langle v_{i(1)}, \ldots, v_{i(k)} \rangle = V$.

Theorem 2.2 *The stopping time τ defined above is strong uniform.*

Proof Suppose $\tau = k$ and let $B = \{v_{i(1)}, \ldots, v_{i(k)}\}$. By definition there exist a subset of vectors $U = \{u_1, \ldots, u_n\} \subset B$ such that $\langle u_1, \ldots, u_n \rangle = V$. Observe that $c_1 \cdot u_1 + \cdots + c_n \cdot u_n$ is uniform in V if c_i, $1 \leq i \leq n$ are independent and uniform in $\mathbb{F}_q$. Therefore $X_k = a_1 \cdot v_{i(1)} + \cdots + a_k \cdot v_{i(k)}$ is uniform in V which proves the result. $\square$

Note that in Theorem 2.2 we do not claim that the stopping time τ is perfect. The example $A = V$ shows that τ is not perfect in general. However, there is a large class of geometric random walks for which this is true. We need a few definitions.

Let $A = \{v_1, \ldots, v_m\}$, and let $[m] = \{1, 2, \ldots, m\}$. For every subset $I = \{i_1, \ldots, i_l\} \subset [m]$ define a subspace $L_I = \langle v_{i_1}, \ldots, v_{i_l} \rangle$. Denote $\mathcal{L} = \mathcal{L}(A)$ the lattice of subspaces L_I for all $I \subset [m]$. We say that A is *proper* if there exist a vector $v \in V$ such that $v \notin L_I$ for all $L_I \neq V$.

Theorem 2.3 *Let $A \subset V$ be a proper set of m vectors. Then the strong uniform time τ defined above is perfect.*

Proof Since A is proper there exists $w \in V$ such that $w \notin L_I$ for all $L_I \neq V$. We claim that w is a halting element. Indeed, if $X_t = w$ then $w \in \langle v_{i(1)}, \ldots, v_{i(t)} \rangle = V$ and by construction $\tau = t$. Therefore whenever gets to w it stops there, i.e. w is a halting element. Thus τ is a perfect time. $\square$

Now observe that the Theorem 2.2 can be stated in purely combinatorial terms. By this we mean that we used only those information about A which can be described by linear relations between vectors. Therefore Theorem 2.3 implies that assuming A is proper, we can compute exactly the separation distance ξ using only combinatorial information about A.

Theorem 2.4 *Let $A \in V$ be a proper set of m vectors, let $\mathcal{W}(A)$ be the corresponding geometric random walk, and let $\mathcal{L} = \mathcal{L}(A)$ be the lattice of subspaces. Then the separation series $\xi(z)$ for the random walk $\mathcal{W}(A)$ is given by the formula*

$$\xi(z) = \sum_{L \in \mathcal{L},\, L \neq V} \frac{(-1)^{n-\dim(L)+1}}{1 - j(L)\, z}$$

where $n = dim(V)$, $j(L) = |A \cap L|/m$.

Proof Let $L \in \mathcal{L}(A)$ be a vector subspace. Clearly,

$$P(\langle v_{i_1}, \ldots, v_{i_k} \rangle \subset L) = j(L)^k$$

By the inclusion-exclusion principle,

$$P(\langle v_{i_1}, \ldots, v_{i_k} \rangle \neq V) = \sum_{L \in \mathcal{L},\, L \neq V} (-1)^{n-\dim(L)+1} j(L)^k$$

By Theorem 2.3, the stopping time τ is perfect. We have

$$\xi(z) = \sum_{k \geq 0} P(\tau > k) \cdot z^k = \sum_{k \geq 0} P(\langle v_{i_1}, \ldots, v_{i_k} \rangle \neq V) \cdot z^k$$

$$= \sum_{k \geq 0} \sum_{L \in \mathcal{L},\, L \neq V} (-1)^{n-\dim(L)+1} j(L)^k z^k = \sum_{L \in \mathcal{L},\, L \neq V} \frac{(-1)^{n-\dim(L)+1}}{1 - j(L)\, z}$$

which proves the claim. □

From Theorem 2.4 one can immediately deduce various properties of the random walk $\mathcal{W}(A)$. In particular, one can obtain the second largest eigenvalue, which can be interpreted as a radius of convergence ρ of the separation series $\xi(z)$ (see [P1]).

Corollary 2.5 *Let A, $\mathcal{W}(A)$, and $\mathcal{L}(A)$ be as in Theorem* 2.4. *Then*

$$s(k) \sim C \cdot \rho^k$$

where $s(k)$ is the separation distance for the random walk $\mathcal{W}(A)$, and

$$\rho = \max_{L \in \mathcal{L}(A)} j(L)\,, \quad C = |\{L \in \mathcal{L}(A), j(L) = \rho\}|\,.$$

Proof. Clear. □

Before we move to particular cases, let us point out to the following straightforward generalization of the results in this section.

Let Q be any set of subspaces of the vector space $V \simeq \mathbb{F}_q^d$. Assume that the vector spaces in Q generate V. Let $\mathbf{P}$ be a probability distribution on Q. Consider a Markov chain X_t on V such that $X_0 = O$ and

$$X_{t+1} = X_t + v$$

where $v = v(t)$ is a vector chosen uniformly randomly from the subspace $L(t) \in Q$, and the subspace $L(t)$ was sampled from Q according to the probability distribution $\mathbf{P}$. Denote this Markov chain by $\mathcal{W}(Q, \mathbf{P})$. Clearly, when Q is a set of lines and $\mathbf{P}$ is uniform, $\mathcal{W}(Q, \mathbf{P})$ is a geometric random walk.

Define the stopping time τ in this case to be the first time we get $L(1) + \cdots + L(\tau) = V$. Denote by $\mathcal{L}(Q)$ the lattice of subspaces of V generated by subspaces in Q. We say that Q is *proper* if there exist a vector $v \in V$ such that $v \notin L$ for all

$L \in \mathcal{L}(Q)$, $L \neq V$. Denote by $\mu(L) = \mu(L, V)$ the Möbius function on L (see e.g. [BBR,S]). Since $\mathcal{L}(Q)$ is a modular lattice, we have $\mu(L) = \pm 1$ (see e.g. [BBR,S]).

Theorem 2.6 *Let Q be a proper set of vector subspaces. Then the stopping time τ defined above is strong uniform. Moreover, if Q is proper then τ is perfect and*

$$\xi(z) = \sum_{L \in \mathcal{L}(Q),\, L \neq V} \frac{\mu(L)}{1 - j(L)\, z}$$

where $j(L) = \sum_{L' \in Q, L' \subset L} \mathbf{P}(L)$.

Clearly, this theorem generalizes Theorems $2.2 - 2.4$. The proof is completely analogous and will be omitted. We challenge the reader to find some interesting applications of this formula.

3. The case of a cube

Let $q = 2$, $m = d$, $V \simeq \mathbb{F}_q^m$, and $A = \{\, (0, \ldots, 1_i, \ldots, 0), 1 \leq i \leq m\}$. Then a geometric random walk $\mathcal{W}(A)$ is equivalent to the lazy random walk on a m-dimensional cube (see section 1).

Random walk on a cube is probably one of the oldest and most thoroughly studied problem (see e.g. [DGM,P1] for results and references). None of the results in this section are new. We nevertheless include them here for methodological reasons as the first nontrivial application of the method, as well as for completeness.

Theorem 3.1 *The separation series $\xi(z)$ for the random walk $\mathcal{W}(A)$ is given by the formula*

$$\xi(z) = \sum_{k=1}^{m} \frac{(-1)^{k+1} \binom{m}{k}}{1 - \frac{m-k}{m}\, z}\,.$$

Proof. In this case $\mathcal{L}(A)$ is a Boolean lattice of coordinate subspaces (see e.g. [BBR,S]). Thus the number of subspaces $L \in \mathcal{L}(A)$ of dimension k is equal to $\binom{m}{k}$, and for each such L we have $j(L) = \frac{m-k}{m}$. Also, A is proper since the vector $(1, \ldots, 1) \in V$ does not belong to any coordinate subspaces except V. Together with Theorem 2.4 this implies the result. $\square$

Theorem 3.2 *Let $\xi = \xi(1)$ be the total separation for the random walk $\mathcal{W}(A)$. We have*

$$\xi = m \cdot \mathfrak{h}(m)$$

where $\mathfrak{h}(m) = 1 + \frac{1}{2} + \frac{1}{3} + \cdots + \frac{1}{m}$.

Proof. Recall that τ is perfect in this case. Use the formula $\xi = E(\tau)$. Finding the expectation of τ is the classical coupon collector's problem. Indeed, we check random coordinates one at a time and stop when all coordinates are checked. Adding the expected time to get the first coordinate, second coordinate, etc, we get

$$\xi = E(\tau) = \frac{m}{m} + \frac{m}{m-1} + \cdots + \frac{m}{1}$$

which proves the result. $\square$

The Euler's formula for the asymptotic behavior of $\mathfrak{h}(m)$ gives us

$$\xi = m \cdot \mathfrak{h}(m) = m \ln(m) + \gamma\, m + \frac{1}{2} + O\left(\frac{1}{m}\right)$$

where $\gamma \approx 0.5772156649$ is the Euler-Mascheroni constant (see e.g. [WW]).

4. Random matroid process

Let S be a finite set and $r : 2^S \to \mathbb{Z}_+$ be a *rank function*. We say that a pair $M = (S, r)$ is a *realizable matroid* over the field $\mathbb{F}_q$ if there exist d and a map $\nu : S \to \mathbb{F}_q^d$ which preserves rank function. An image $A = \nu(S)$ is called *realization* of a matroid $M = (S, r)$. By Theorem 2.4, we have the following result.

Proposition 4.1 *If A is a proper set of vectors, then the separation series $\xi(z)$ of the random walk $\mathcal{W}(A)$ depends only on a matroid (S, r) and not on the realization A.*

Proof. Clear. □

It is easy to see that if (S, r) is realizable over $\mathbb{F}_q$ then it is realizable over any $\mathbb{F}_{q'}$, such that $q' > q$ (see e.g. [A]). Thus one can consider realizations over fields with sufficiently large q.

Proposition 4.2 *If M is a realizable matroid over the field $\mathbb{F}_q$, and q is sufficiently large, then every realization $A \subset \mathbb{F}_q^d$ is proper.*

Proof. The maximum number $N(q)$ of points in all $(d-1)$-dimensional subspaces $L \in \mathcal{L}(A)$ of any realization $A \subset \mathbb{F}_q^d$ is bounded by a polynomial of q degree $d-1$. Therefore for sufficiently large q we have $|V| = q^d > N(q)$. This implies the result. □

Now consider the following random process $\mathcal{B} = \mathcal{B}(M)$. Fix a realizable matroid $M = (S, r)$, $r(S) = d$. Let $B_0 = \emptyset$, $B_{t+1} = B_t \cup s$ where $s = s(t) \in S$ is chosen uniformly. Clearly $r(B_t) \le r(B_{t+1})$. Stop the first time t such that $r(B_t) = d$. We call B_t the random matroid processes. Denote by κ the stopping time of the process $\mathcal{B}(M)$. Theorem 2.4 combined with Proposition 4.1, 4.2 gives us the following result.

Theorem 4.3 *Let $M = (S, r)$ be a realizable matroid such that $r(S) = d$. Let $A \subset \mathbb{F}_q^d$ be realization of M, and let κ be the stopping time of the random process $\mathcal{B}(M)$. Consider a geometric random walk $\mathcal{W}(A)$. Then for large enough q we have $\xi = E(\kappa)$.*

Proof. Clear. □

We finish this section by constructing proper realizations of the graphical matroids.

Let Γ be a simple connected graph (no orientation, no loops, no multiple edges) with vertex set Y, and edge set $E \subset Y \times Y$. Consider a rank function $r : 2^E \to \mathbb{Z}_+$ defined as follows:

$$r(H) = |Y| - c(Y, H)$$

where $H \subset E$, and $c(Y,H)$ is the number of connected components of a subgraph (Y,H). By definition, $r(E) = |Y| - 1$. We call (E,r) a *graphical matroid*.

Now, choose any vertex $y_0 \in Y$ to be a *root*. Fix an orientation of the edges towards the root. For any $q \geq 2$ consider the following realization $A = \nu(S) \subset \mathbb{F}_q^{|Y|-1}$ of a matroid (E,r):

$$\nu(y,y_0) = e_y\,, \quad \nu(y,y') = e_y - e_{y'}, y' \neq y_0$$

for all $(y,y_0),(y,y') \in E$, and where e_y, $y \in Y - y_0$ is a basis in $\mathbb{F}_q^{|Y|-1}$.

Proposition 4.4 *For any $q \geq 2$ the set of vectors $A = \nu(S) \subset \mathbb{F}_q^{|Y|-1}$ is a realization of a matroid (S,r). Moreover, if $q \geq |Y|$, this is a proper realization.*

Proof. The first part is well known in the literature (see e.g. [A]). A straightforward check shows that vector $w = \sum_{y \in Y - y_0} e_y$ does not belong to any subspace $L \in \mathcal{L}(A)$, such that $dim(L) < |Y| - 1$. This proves the second part. □

Now consider the following random process. Let $H_0 = \emptyset$, $H_{t+1} = H_t \cup (y_1,y_2)$ where $(y_1,y_2) \in E$ is a edge of graph Γ chosen uniformly. Denote κ the first time t such that subgraph (Y,H_t) is connected. By definition, the random graph process H_t corresponds to a random matroid process for B_t in this case. As before, denote by κ the stopping time of this process. Theorem 4.3 combined with Proposition 4.4 gives us the following result.

Theorem 4.5 *Let Γ be a simple graph with n vertices, (S,r) be the corresponding graphical matroid, and $A = \nu(S)$ its realization over $\mathbb{F}_q$, $q \geq n$. Consider a geometric random walk $\mathcal{W}(A)$ and its total separation distance ξ. We have $\xi = E(\kappa)$.*

Proof. Clear. □

Remark 4.6 Note that the random graph process we consider is somewhat different from the random graph process normally studied in random graph theory (see [Bo]).

5. The case of a complete graph.

Suppose A contains vectors e_l, $1 \leq l \leq n-1$, and $e_i - e_j$, $1 \leq i < j \leq n-1$, where $e_1, \ldots, e_{n-1}$ is a basis in $V \simeq \mathbb{F}_q^{n-1}$. It is easy to see that A is a realization of a graphical matroid which corresponds to the complete graph $\Gamma = K_n$. We have $Y = [n]$, $|E| = \binom{n}{2}$.

Again, this is a well understood case, which we include here for completeness and as a first application of the approach.

Theorem 5.1 *Let A be as above, $\mathcal{W}(A)$ be the corresponding random walk, and ξ be its total separation distance. Then*

$$C_1 n \log n \leq \xi \leq C_2 n \log n$$

for some absolute constants C_1, C_2.

The result is clearly not sharp. Later, in Example 8.4, we will prove the cutoff phenomenon in this case. The following proof, however, is a starting point of several generalizations.

Proof By Theorem 4.5 the total separation distance ξ for this random walk is equal to the expected time to get a connected subgraph when adding one random edge at a time. We shall bound this expectation from above and below.

Let (i_1, i_2) be the first edge. Denote H_t the set of edges after t steps, and $C_t \subset Y$ the component of the graph (Y, H_t) which contains i_1. Let $k = |C_t|$ be the number of vertices in C_t. Denote p_k the probability that $|C_{t+1}| > |C_t|$. Since the total number of edges between C_t and $[n] \setminus C_t$ is $k \cdot (n-k)$, we have

$$p_k \geq \frac{k \cdot (n-k)}{\binom{n}{2}}$$

Let κ be the first time $|C_t| = n$. We have

$$E(\kappa) \leq 1 + \sum_{k=2}^{n-1} \frac{1}{p_k} = 1 + \binom{n}{2} \cdot \sum_{k=2}^{n-1} \frac{1}{k \cdot (n-k)}$$

$$= 1 + \frac{n(n-1)}{2} \cdot \frac{1}{n} \cdot \sum_{k=2}^{n-1} \frac{1}{k} + \frac{1}{n-k} \sim n \log n$$

This gives the upper bound. To prove the lower bound, denote ι the first time each vertex in a subgraph (Y, H_t) has at least one adjacent edge. Think of the edges we add as of a random pair of elements in $Y = [n]$. Since at a time we cannot do better than add two vertices, this becomes a coupon collector's problem again and $E(\iota) = \Omega(n \log n)$. Observe that $E(\iota) \leq E(\kappa)$, which implies the lower bound. □

6. The case of coordinate subspaces.

Let $e_1, \ldots, e_n$ be a basis in $V \simeq \mathbb{F}_q^n$. For every $I = \{i_1, \ldots, i_k\} \subset [n]$ consider a *coordinate subspace* $L_I = \langle e_{i_1}, \ldots, e_{i_k} \rangle$. Denote by $Q(n,k)$ the set of all k-dimensional coordinate subspaces. Clearly, $|Q(n,k)| = \binom{n}{k}$.

Theorem 6.1 *Let $Q = Q(n,k)$ be as above, $\mathcal{W}(Q)$ be the corresponding random walk, and ξ be its total separation distance. Then*

$$\xi = \Theta\left(\frac{\log n}{\log n - \log(n-k)}\right).$$

Proof. Consider the following variation on the coupon collector's problem. Instead of taking 1 coupon at a time, now take a random sample of k different coupons at a time. Denote by T the time all n coupons are collected. The problem is to find $E(T)$. Let us show how one can imbed this problem in the usual coupon collector's problem.

Denote by E_l the expected time to collect l different coupons in the usual coupon collector's problem. Recall that

$$E_l = \frac{n}{n} + \frac{n}{n-1} + \cdots + \frac{n}{n-l-1} = n\left(\mathfrak{h}(n) - \mathfrak{h}(n-l)\right).$$

Therefore the expected time E' to get all the coupons if given k different coupons at once is given by

$$E' \sim \frac{E_n}{E_k} = \frac{\mathfrak{h}(n)}{\mathfrak{h}(n) - \mathfrak{h}(n-l)}.$$

Observe that by Theorem 2.6 we have $\xi = E(T)$. This proves the result. □

7. The case of vectors in generic position.

One of the interesting recently studied questions concerns the behavior of the *random random walks* (see e.g. [DH,P4,R]). These are basically random walks on a fixed group with a set of generators randomly chosen from a given distribution. In this section we will study random geometric random walks which as we show correspond to the case of lines in generic position.

Let A be a set of n vectors in $V \simeq \mathbb{F}_q^k$. We say that A is *generic* if every k vectors in A are linearly independent.

Theorem 7.1 *Let A be a set of n vectors in $V \simeq \mathbb{F}_q^k$. Let $\mathcal{W}(A)$ be the corresponding geometric random walk, and ξ be its total separation distance. Then*

$$\xi \geq n \cdot (\mathfrak{h}(n) - \mathfrak{h}(n-k))$$

and the equality holds if and only if A is generic.

Proof For any A, we need at least k different vectors to generate V. By coupon collector's problem, the expected time to get these k vectors is $E = n \cdot (\mathfrak{h}(n) - \mathfrak{h}(n-k))$. Now, by Theorem 4.3 we immediately have $\xi = E$, when A is generic and $\xi > E$ otherwise. $\square$

8. The cutoff phenomenon

Let (G_i), (S_i), $i = 1, 2, \ldots$ be a sequence of groups and generating sets. Consider a sequence of random walks $(\mathcal{W}_i)$. Denote by $s_i(\cdot)$ and ξ_i the corresponding separation distance and the total separation.

We say that a sequence of random walks $(\mathcal{W}_i)$, $i = 1, 2, \ldots$ has a *cutoff* if there exist two integer sequences (a_i) and (b_i) such that $a_i/b_i \to 1$, $s_i(a_i) \to 0$ and $s_i(b_i) \to 1$ as $i \to \infty$. This definition is due to Aldous and Diaconis (see [AD,D2]).

Example 8.1 Suppose $G = \mathbb{Z}_2^m$ and $\mathcal{W}$ is a random walk on a cube (see section 3). Recall that the perfect stopping time τ is defined a time to check all the coordinates. We have

$$\xi = E(\tau) = m \cdot \mathfrak{h}(m) = m \log m + O(m)$$

where $\mathfrak{h}(m) = 1 + \frac{1}{2} + \frac{1}{3} + \cdots + \frac{1}{m}$ Recall that the element $v = (1, 1, \ldots, 1)$ is halting and thus $s(k) = 1 - 2^n P^k(v) = P(\tau \leq k)$. Now, a direct computation for coupon collector's problem shows that

$$Var(\tau) = m \sum_{i=1}^{m-1} \frac{i}{(m-i)^2}$$

(see e.g. [F, §9.9]). From here we have

$$Var(\tau) \leq m^2 \sum_{i=1}^{\infty} \frac{1}{i^2} = \frac{\pi^2}{6}\, m^2$$

and by Chebyshev inequality

$$s(m \cdot \mathfrak{h}(m) - x \cdot m) \leq \frac{C_1}{x^2}$$

$$s(m \cdot \mathfrak{h}(m) + x \cdot m) \geq 1 - \frac{C_1}{x^2}$$

for some absolute constant C_1. This shows cutoff for the random walk on a cube of dimension m.

Now we can generalize this observation. Let $V = \mathbb{F}_q^n$, let $A \in V$ be a set of vectors, and let $\mathcal{W} = \mathcal{W}(A)$ be the geometric random walk. Consider the corresponding matroid M and the random matroid process $\mathcal{B}$. Also, let κ be the stopping time of $\mathcal{B}$.

We say that a sequence of random matroid processes $(\mathcal{B}_i)$ has a *sharp threshold* if there exist two integer sequences (a_i) and (b_i) such that $a_i/b_i \to 1$, $P(\kappa_i > a_i) \to 0$ and $P(\kappa_i < b_i) \to 0$ as $i \to \infty$.

Theorem 8.2 *Let $A_i \in V_i$, $i = 1, 2, \ldots$ be proper sets of vectors. Then the sequence of random walks $(\mathcal{W}_i)$ has a cutoff if and only if $(\mathcal{B}_i)$ has a sharp threshold.*

Proof. Let $v_i \in V_i$ be a halting vector, and let τ_i be the corresponding perfect time. For every i we have

$$s_i(k) = 1 - P^k(v_i) \cdot |V_i| = P(\tau_i \leq k) = P(\kappa_i \leq k)\,.$$

This implies the result. □

Proposition 8.3 *If $Var(\kappa_i)/E(\kappa_i)^2 \to 0$ $E(\kappa) \to \infty$ as $i \to \infty$, then $(\mathcal{B}_i)$ has a sharp threshold.*

Proof. Take

$$a_i = E(\kappa_i) + w(i) \cdot \sqrt{Var(\kappa_i)}$$

$$b_i = E(\kappa_i) - w(i) \cdot \sqrt{Var(\kappa_i)}$$

where $w(i)$ is an increasing function

$$w(i) = \left(\frac{E(\kappa_i)^2}{Var(\kappa_i)} \right)^{\frac{1}{4}} \to \infty \quad \text{as } i \to \infty.$$

Now use Chebyshev inequality to show sharp threshold. □

Example 8.4 *(The case of a complete graph.)*
Let A_n be a proper realization of a graphical matroid which corresponds to complete graph $\Gamma = K_n$ (see section 5). Consider the corresponding random walk $\mathcal{W}_n = \mathcal{W}(A_n)$. We shall prove that $\xi = \frac{1}{2} n \log n + O(n)$ and that in fact we have a cutoff in this case.

Indeed, consider the corresponding random graph process $\mathcal{B}_n$. We take an empty graph and keep adding random edges until the obtained subgraph of K_n is connected. Let κ_n be the corresponding stopping time. By Theorem 8.2, we need to

show that $(\mathcal{B}_n)$ has a sharp threshold. But this is a known result in the theory of random graphs.

Consider a random graph process $\mathcal{B}'_n$ which works in a similar way but now let us not allow repetition of edges. In other words, each time we choose an edge which is a random edge which is not in our graph. The corresponding stopping time κ_n^{pr} will always be bounded by $\binom{n}{2}$, which is the total number of edges in K_n.

Now, for the random processes $(\mathcal{B}'_n)$ Erdős and Rényi showed a very sharp threshold. Namely, they showed that for $k = (n/2)(\log n + x + o(1))$

$$P(\kappa' \leq k) \to e^{-e^{-x}}$$

(see [ER; Bo, §9.1]).

To apply this result in our situation, observe that among m edges chosen randomly we get on average $O(m^2/n^2)$ repetitions. Indeed, when a new edge is added to our set of j different edges it is one of the previous edges with probability $j/\binom{n}{2}$. Thus the average number of repetitions is at most $\sum_{j=1}^{m} j/\binom{n}{2} = O(m^2/n^2)$. In our case $m = O(n \log n)$ which gives us $O(\log^2 n)$ repetitions. Since this is small compared to $O(n \log n)$, the Chernoff bound implies that the probability to get more than $O(\log^2 n)$ number of repetitions is exponentially small. From here we have $E(\kappa_n) = 1/2\, n \log n + O(n)$ and $(\mathcal{B}_n)$ has a sharp threshold. Therefore, $(\mathcal{W}_n)$ has a cutoff.

Remark 8.5 The cutoff considered in this paper is different from the cutoff considered in [D2] and other papers, where the variation distance $tv(k)$ was considered instead of the separation distance. While it is similar in flavor, it is not clear to us whether either of them implies the other. Still, preliminary computations seem to indicate that cutoff for the total variation distance is a somewhat stronger condition.

9. Weak threshold of Margulis.

In this section we recall a pioneer result of Margulis (see [M]), who showed that a slightly weaker version of "sharp threshold" exists for *all* matroid sequences.

Let M be a vector matroid. Denote $\eta(M)$ the smallest number of vectors to be removed to obtain a matroid of a smaller rank. In case of a graphical matroid, this is simply the size of the minimal cut. Also, denote by $|M|$ the number of vectors in M.

Let (M_i), $i = 1, 2, \ldots$ be a sequence of matroids, and let $\mathcal{B}'_i = \mathcal{B}'(M_i)$ be the corresponding sequence of random matroid processes where repetitions are not allowed. Also, let $\kappa'_i = \kappa'_i(\mathcal{B}_i)$ be their stopping times. For every ϵ denote by $a_i(\epsilon)$ and $b_i(\epsilon)$ the smallest and the largest numbers respectively such that

$$P(\kappa_i \leq a_i) \geq 1 - \epsilon$$
$$P(\kappa_i \leq b_i) \leq \epsilon$$

Theorem 9.1 (Margulis) *Suppose $\eta(M_i) \to \infty$ as $i \to \infty$. Then:*

$$\frac{a_i - b_i}{|M_i|} \to 0 \ , \quad \text{as } i \to \infty$$

While the result of Margulis is similar to what we need to establish a cutoff, it gives a bit weaker bounds. Indeed, in case of a complete graph K_n we get $a_n - b_n = o(n^2)$ while to establish a sharp threshold and a cutoff for the corresponding geometric random walk we need $a_n - b_n = o(n \log n)$, which we know is true. In the next example we show that one can successfully apply Theorem 9.1 to show cutoff.

Example 9.3 Consider a *"thick line"* graph C_n defined by connecting vertices i and $i+1$ by $\log n$ edges, $i = 1, 2, \ldots, n$. Consider a random graph process. Since $\log n \to \infty$ the repetitions are rare, so when n is large both processes $\mathcal{B}(C_n)$ and $\mathcal{B}'(C_n)$ have about the same stopping time. To get a connected spanning subgraph we need an edge between every pair of vertices i and $i+1$. Thus by coupon collector's problem $E(\kappa) = n \log n + O(n)$. On the other hand, $\eta(C_n) = \log n \to \infty$ and Margulis' theorem implies that $a_i - b_i = o(|C_n|) = o(n \log n)$ and thus we have a cutoff.

This example can be easily generalized to the case of vectors in generic position (see section 7).

10. The case of random set of vectors.

For a given n, q and m consider a random set of vectors $A \subset V = \mathbb{F}_q^n$, $|A| = m$. What can we say about the total separation ξ of the geometric random walk $\mathcal{W}(A)$. Would there be a cutoff?

Clearly, $n \le m$ or otherwise vectors in A will not generate V. Suppose $l = m - n$ is fixed and n grows. Roughly, we need to use almost all the vectors to generate the whole space V. Thus by coupon collector's problem we need about $n \log n$ walk steps. This can be formalized by the following result.

We call a m-tuple a set of vectors $A \subset V$, such that $|A| = m$.

Theorem 10.1 *For any $\epsilon, \delta > 0$ there exist constants c_1, c_2, l_1, n_1 such that for every $n \ge n_1$, $l \ge l_1$ a random $(n+l)$-tuple A satisfies the following inequalities*

$$s(n \log n + c_1\, n) \le \epsilon$$

$$s(n \log n + c_2\, n) \ge 1 - \epsilon$$

with probability $1 - \delta$, where $s(k)$ is the separation distance after k steps of the corresponding random walk $\mathcal{W}(A)$.

This roughly means that as $n \to \infty$, a sequence of random $(n+l)$-tuples has a cutoff. Heuristically, this implies that for almost all sets A the mixing time is about the fastest possible. Thus we have a cutoff.

To prove the theorem we need several observations. We say that set $B \subset V$ is *full* if the linear span of B generates the whole space $V = \mathbb{F}_q^n$. Suppose $|B| = m$, $m \ge n$. Compute the probability $P(n, m)$ that a random m-tuple B is full. Combine vectors in B into a $m \times n$ matrix M. We have

$$P(n,m) = P(rk(M) = n) = \left(1 - \frac{1}{q^m}\right) \cdot \left(1 - \frac{1}{q^{m-1}}\right) \cdot \ldots \cdot \left(1 - \frac{1}{q^{m-n+1}}\right)$$

Thus the probability that a random $(n+l)$-tuple is not full is roughly $1-q^{1-l}$. Let l_0 be the smallest integer such that for large enough n and $l \geq l_0$ we have $P(n, n+l) > 1 - \epsilon\,\delta$.

Now, let $l_1 = l_0 + 1$. An r-subtuple B in m-tuple A is an r-tuple which is a subset of A: $B \subset A$.

Lemma 10.2 *Suppose $l \geq l_1$. Then at least $1-\delta$ fraction of all $(n+l)$-tuples A satisfies the following property:*

($\circ$) *A random $(n+l_0)$-subtuple B of A (chosen uniformly of all $\binom{n+l}{n+l_0}$ subtuples) is full with probability at least $1-\epsilon$.*

Proof. Let $A_1, \ldots, A_N$ be the set of all $(n+l)$-tuples of V. For each A_i let p_i be the probability that a random $(n+l_0)$-tuple of A_i is not full. Also, let p be the probability that a random $(n+l)$-tuple violates the property ($\circ$) described in the Lemma. Notice that by definition:

$$\frac{1}{N}\sum_{i=1}^{N} p_i \geq \epsilon\, p.$$

On the other hand, by the symmetry, $\sum_{i=1}^{N} p_i$ is N times the probability that a random $(n+l_0)$-tuple chosen from V is not full. The last probability is at most $\epsilon\,\delta$. Therefore

$$\epsilon\,\delta \geq \epsilon\, p\,,$$

which implies $p \leq \delta$, completing the proof. □

Proof of Theorem 10.1. From above, with probability at least $1-\delta$ almost all ($(1-\epsilon)$ fraction to be exact) of $(n+l_0)$-subtuples in a random $(n+l)$-tuple are full. This implies that with probability at least $1-\delta$ the expected stopping time of a random matroid process $\mathcal{B}(A)$ is $n\log n + O(n)$. This immediately implies desired result about mixing and finishes proof of the theorem. □

Remark 10.3 There are various other results about the behavior of the so called *random random walks* (see e.g. [DH,P4,R]) and the connection with cutoff phenomenon (see [D2] for review and references). Notably, in papers [G,W] the cutoff in terms of variation distance was shown for almost all sets of generators of $\mathbb{Z}_2^n$. While this result roughly corresponds to the case $q=2$, the technique is different from ours.

11. The case of random graphs.

By analogy with the previous section one can consider a sharp threshold of random graph process for graphs with n vertices and m edges. It turns out that one can prove results similar to those of the previous section. The applications to the cutoff of geometric random walks are clear, so we will deal directly with graphs.

Let G be a graph on n vertices. We say that G is *m-graph* if $|G| = m$, i.e. G has m edges. We say that set H is *l-subgraph* if $H \subset G$ and $|H| = l$.

It is known that for $m = o(n\log n)$ almost every random m-graph is disconnected. We will show that in a sense $\frac{1}{2}n\log n$ edges is enough not only for a graph to be connected, but also to have sharp threshold.

Theorem 11.1 *For any $\epsilon, \delta > 0$ there exist constants c_1, c_2, l_1, n_1 such that for every $n \geq n_1$, $l \geq l_1$, a random $\left(\frac{1}{2} n \log n + l\, n\right)$-graph G on n vertices satisfies the following inequalities*

$$P(\kappa \leq n \log n + c_1\, n) \leq \epsilon$$

$$P(\kappa \leq n \log n + c_2\, n) \geq 1 - \epsilon$$

with probability $\geq 1 - \delta$, where κ is the stopping time of the random graph process $\mathcal{B}(G)$.

Proof of this theorem is largely the same as of Theorem 10.1, so we will simply sketch it.

First, we can use the bounds cited in Example 8.4 to find l_0 such that a random $\left(\frac{1}{2} n \log n + l_0\, n\right)$-graph G on n vertices is connected with probability at least $1 - \epsilon\, \delta$. Take $l_1 = l_0 + 1$.

Now, refer to the following analog of Lemma 10.2:

Lemma 11.2 *Suppose $l \geq l_1$. Then at least $1 - \delta$ fraction of all $\left(\frac{1}{2} n \log n + l\, n\right)$-graph G satisfies the following property:*

($\circ$) A random $\left(\frac{1}{2} n \log n + l_0\, n\right)$-subgraph H of G (chosen uniformly of all the subgraps) contains a spanning tree with probability at least $1 - \epsilon$.

Again, the lemma applied to the general set immediately proves the theorem. □

12. The case of edge-transitive graphs.

In this section we show that if we have a sequence of edge-transitive graphs, it always has a sharp threshold. Thus we effectively "derandomize" the result of the previous section.

Graph G is called *edge-transitive* if for every pair of edges E_1 and E_2 there is an automorphism $\pi : G \to G$ such that $\pi(E_1) = E_2$. For example, both complete graph K_n and complete bipartite graph $K_{m,n}$ are edge-transitive.

Let $d = d(G)$ denote the minimum degree of G.

Theorem 12.1 *Let (G_n) be a sequence of edge-transitive graphs on n vertices such that $\log d / \log n \to 0$ as $n \to \infty$. Then a random graph process $\mathcal{B}'(G_i)$ has a sharp threshold.*

We prove Theorem 12.1 using another formulation, more traditional for random graphs.

Suppose we choose each edge of G independently with probability p, and let G_p denote the resulting random graph. Suppose p_0 and p_1 satisfy $P(G_{p_0} \text{ is connected}) = \epsilon$ and $P(G_{p_1} \text{ is connected}) = 1 - \epsilon$. We need to show that $u(G, \epsilon) = (p_1 - p_0)/p_0 \to 0$ as $n \to \infty$, where G is as in the theorem.

Let us formalize the problem again. Suppose a graph G has n vertices and m edges. Consider the hypercube $\mathbf{H} = \{0, 1\}^m$, where each coordinate corresponds to an edge in G (so each vertex of the cube corresponds to a subgraph of G). Here we will identify sets of subgraphs of G with sets of vertices of the hypercube. Define a probabilistic space on H by setting $\mu_p(x) = p^{w(x)}(1-p)^{m - w(x)}$ for any vertex x, where $w(x)$ is the number of 1's in x. Furthermore, for each $x \in \mathbf{H}$, let x^i denote the vertex of $\mathbf{H}$ obtained by changing the i^{th} coordinate of x.

Let A be a monotone subset of $\mathbf{H}$, and $A(p) = \mu_p(A) = \sum_{x \in A} \mu_p(x)$. For each coordinate i define $A_i = \{x | x \in A, x^i \notin A\}$, and set $\gamma(p) = max_i A_i(p)$.

The core of the proof is the following lemmas, shown by Talagrand [T] and Margulis [M], respectively.

Lemma 12.2 (Talagrand) *There is a constant $c > 0$ such that:*

$$p\frac{\partial A(p)}{\partial p} \geq c\frac{\log(1/\gamma(p))}{\log(1/p)} A(p)(1 - A(p)).$$

Lemma 12.3 (Margulis)

$$p\frac{\partial A(p)}{\partial p} = 2\sum_{i=1}^{m} A_i(p).$$

Now let A be the set of all connected spanning subgraphs of G. Fix an arbitrarily large C. Assume that $u(G, \epsilon) < C$. It follows by classical analysis that there is a point $p \in [p_0, p_1]$ such that $p\frac{\partial A(p)}{\partial p} < C$.

Since G is edge-transitive, A is symmetric, so $A_i(p)$ are equal. Therefore, Margulis lemma implies that

$$\gamma(p) = \frac{1}{m}\sum_{i=1}^{m} A_i(p) < \frac{C}{m}$$

Combine this with Talagrand's lemma, we have:

$$C' = \frac{C}{c\epsilon(1-\epsilon)} \geq \frac{\log(m/C)}{\log(1/p)}$$

which yields

$$\log(1/p) > C'' \log n$$

for some constant C'' depending on C (here we use the fact that $m \geq n/2$).

Recall that the minimum degree of G is d, so the probability that a v with degree d is isolated is $(1-p)^d$. Since $A(p) \geq \epsilon$, it follows that $(1-p)^d \leq 1 - \epsilon$, therefore $p \geq c/d$ for some constant $c = c(\epsilon)$. Thus

$$\log(1/p) < c' \log d$$

for some constant c'. This leads to a contradiction, since $\log d = o(\log n)$. $\square$

The theorem covers many nice symmetric cases such as cycle, thick cycle (see Example 9.3), hypercubes (see Example 8.1) and many others.

Example 12.4 Let $0 < \alpha < 1$ and $\begin{bmatrix} n \\ \alpha n \end{bmatrix}$ be the set of all $\lfloor \alpha n \rfloor$-subsets of $[n] = \{1, \ldots, n\}$. Consider a sequence of graphs Υ_n with vertices in $\begin{bmatrix} n \\ \alpha n \end{bmatrix}$ and edges (I, J), $I, J \in \begin{bmatrix} n \\ \alpha n \end{bmatrix}$, such that $|I \cap J| = \lfloor \alpha n \rfloor - 1$. Clearly, Υ_n is edge-transitive. It

is easy to see that $d(\Upsilon_n) = n - \lfloor \alpha n \rfloor = O(n)$ while $\left| \begin{bmatrix} n \\ \alpha n \end{bmatrix} \right| = \binom{n}{\lfloor \alpha n \rfloor} = O(\beta^n/\sqrt{n})$, where $\beta = (\alpha^\alpha (1-\alpha)^{1-\alpha})^{-1} > 1$. Thus $\log d / \log n \to 0$ as $n \to \infty$ and by Theorem 12.1 the sequence of random graph processes $\mathcal{B}'(\Upsilon_n)$ has a sharp threshold.

Example 12.5 Let $\Gamma = K_{m,n}$ be a complete bipartite graph, $m \geq n$. We have $d = m$ and Γ is edge-transitive with $m + n$ vertices. Theorem 12.1 implies that the sequence of the corresponding random graph processes $\mathcal{B}'(K_{m(n),n})$ has a cutoff if $\log n / \log m(n) \to 0$ as $n \to \infty$. For example, $m(n) = n^{\log n}$ will work.

Example 12.6 Let Γ_n be a sequence of Cayley graphs of finite groups generated by small conjugacy classes. Then Γ_n are all edge-transitive graphs and the theorem implies that we have a sharp threshold for the corresponding random graph processes.

For example, let $G = S_n$ be a symmetric group, and T be a set of all transpositions. Let Γ_n be the corresponding Cayley graph. It is clearly edge-transitive. We have

$$\frac{d(\Gamma_n)}{|\Gamma_n|} = \frac{|T|}{|S_n|} = \frac{\binom{n}{2}}{n!} \to 0 \quad \text{as } n \to \infty$$

Now use Theorem 12.1 to show the sharp threshold.

Theorem 12.1 can be generalized to transitive matroids. We call matroid $M = (S, r)$ *transitive* if for any two elements $s_1, s_2 \in S$ there exist a permutation $\pi : S \to S$ such that $\pi(s_1) = s_2$ and $r(\pi(X)) = r(X)$ for every $X \subset S$. A matroid of the form $M' = (S', r|_{S'})$, $S' \subset S$ is called *submatroid*. A matroid (S, r) is called *connected* if it is not a sum of two sumbmatroids $(S', r') + (S'', r'')$, where $S = S' \cup S''$, $S' \cap S'' = \emptyset$, and $r(X) = r'(X \cap S') + r''(X \cap S'')$ for all $X \subset S$.

The role of vertices (or rather their complements) of a matroid M is now played by connected submatroids M' such that $r(M') = r(M) - 1$. We call these *generalized vertices*. Let the *degree* of such a generalized vertex be the number of elements $s \in S$ such that $r(M' \cap \{s\}) = r(M)$. Define the *degree* of a matroid to be the minimum degrees of its generalized vertices.

Theorem 12.7 *Let (M_n) be a sequence of transitive matroids with n generalized vertices and degree $d = d(n)$ such that $\log d(n) / \log n \to 0$ as $n \to \infty$. Then a random matroid process $\mathcal{B}'(M_n)$ has a sharp threshold.*

The proof of Theorem 12.7 is identical to the proof of Theorem 12.1. We use the fact that Margulis and Talagrand lemmas are known for matroids (see [M,T]). We omit the details.

Note that Theorems 12.1 and 12.7 are false if the (crucial) transitivity assumption is dropped. For instance, consider two copies of the hypercube graph of dimension d. In the first copy, delete an edge uv (say), and in the second copy delete an edge $u'v'$. Draw two new edges uu' and vv'. The resulting graph G is connected, d-regular and has 2^{d+1} vertices. But $u(G, \epsilon) \sim (1 - 2\epsilon)/\epsilon$ for all small ϵ.

When the degree d is large, a similar result can be proved under somewhat different assumption. We say that two subgraphs G_1 and G_2 of G are equivalent if there is an element $\pi \in Aut(G)$ such that $G_1 = \pi(G_2)$. For a subforest F, let $\varrho(F)$ be the number of subforests of G equivalent to F. Let $\varrho(s) = \min_{F, |V(F)=s|} \varrho(F)$. Let $w(n)$ be an arbitrary function tending to infinity.

Theorem 12.8 *There is a function $f(C)$ such that for any $C > 0$, there is $n_0 = n(C)$ such that if G is a connected graph on $n > n_0$ vertices and $\varrho(f(C)) > (w(n)d(G))^{f(C)}$ then $u(G, \epsilon) > 1/C$.*

Theorem 12.8, in a way, asserts that if G has a sufficiently large automorphism group, then G has sharp threshold. This theorem covers basis cases such as complete graphs K_n (cf. Example 8.4) or complete bipartite graphs $K_{n,n}$. The proof is more involved and omitted. Theorem 12.8 can be generalized for matroids as well.

13. Odds and ends.

In this paper we showed that the behavior of the geometric random walks is largely determined by random matroid processes. Still, there are many open questions regarding the matter.

1) Is there an analog of the random matroid processes for other types of random walks? The generalization to all abelian groups is straightforward, while the case of nilpotent groups is less obvious. We discuss this in [AP].

Probably, the most interesting case is the random walk on a symmetric group S_n generated by transpositions. In this case there are several constructions of the strong uniform times (see [D1,P1,P2]). The mixing time $\xi = \frac{1}{2} n \log n$ and there are some strong connections with the random process for a complete graph K_n. We discuss these connections at length in [P3].

2) In this paper we were able to establish the existence of a cutoff in some cases, even when we don't know the mixing time. There are other instances (cf. [D2,P1]) when such a result might be desirable.

3) There are still many gaps in our understanding of the sharp threshold phenomenon. For example, we still don't know if there exist a sharp threshold for random graphs with small number of edges (conditioned they are connected).

Also, it is worth noting that the result of Margulis cannot be strengthened to show the cutoff. Suppose $\eta(M_i) \to \infty$ as $i \to \infty$. Then sequence of $\mathcal{B}(M_i)$ may not have a sharp threshold. Indeed, consider two copies of a complete graph K_{n^2} connected by n edges. Consider the corresponding graphical matroid M_n. Then $\eta(M_n) = n \to \infty$, while easy computations show that there is no sharp threshold.

4) After the paper was completed we learned about the result of Friedgut and Kalai [FK] which deals with threshold for symmetric properties, in particular graph properties. As connectivity of random subgraphs is a natural example of such property, which suggests that our results for edge transitive graphs are related to this work. We challenge the reader to uncover this curious connection.

5) The assumption of Theorem 12.1 may be too special or hard to check (although the theorem covers several important graphs); it is natural to ask whether the statement still holds under a more general assumption.

References

[A] M. Aigner, *Combinatorial Theory*, Springer, Berlin, 1979.

[AD] D. Aldous, P. Diaconis, *Strong uniform times and finite random walks*, Advances in Applied Math. **8** (1987), 69–97.

[AF] D. Aldous, J. Fill, *Reversible Markov Chains and Random Walks on Graphs*, monograph in preparation, 1996.

[ASE] N. Alon, J. H. Spencer, P. Erdős, *The Probabilistic Method*, Wiley, New York, 1992.

[AP] A. Astashkevich, I. Pak, *Random walks on nilpotent and supersolvable groups*, preprint (1997).

[Ba] L. Babai, *Automorphism groups, isomorphism, reconstruction*, in Handbook of Combinatorics (R. L. Graham, M. Groetschel, and L. Lovasz, eds.) (1996), Elsevier.

[BBR] M.Barnabei, A. Brini, G-C. Rota, *The theory of Möbius functions,* (in Russian), Uspekhi Mat. Nauk **41** (1986), 113–157.

[Bo] B. Bollobás, *Random graphs*, Academic Press, London, 1985.

[D1] P. Diaconis, *Group Representations in Probability and Statistics*, IMS, Hayward, California, 1988.

[D2] P. Diaconis, *The cutoff phenomenon in finite Markov chains*, Proc. Nat. Acad. Sci. U.S.A. **93** (1996), 1659–1664.

[DF] P. Diaconis, J. A. Fill, *Strong stationary times via new form of duality*, The Annals of Probability **18** (1990), 1483–1522.

[DGM] P. Diaconis, R. Graham, J. Morrison, *Asymptotic analysis of a random walk on a hypercube with many dimensions*, Random Structures and Algorithms **1** (1990), 51–72.

[DH] C. Dou, M. Hildebrand, *Enumeration and random random walks on finite groups*, Ann. Probab. **24** (1996), 987–1000.

[ER] Erdős and Rényi, *On random graphs. I.*, Publ. Math. Debrecen **6** (1959), 290–297.

[F] W. Feller, *An introduction to Probability theory and its applications* (third edition), John Wiley, New York, 1970.

[FK] E. Friedgut, G. Kalai, *Every monotone graph property has a sharp threshold*, Proc. Amer. Math. Soc. **124** (1996), 2993–3002.

[G] A. Greenhalgh, *A model for random random-walks on finite groups*, Combin. Probab. Comput. **6** (1997), 49–56.

[LW] L. Lovász, P. Winkler, *Mixing Times* (1998), AMS DIMACS Series, vol. 41, 189–204.

[M] G. Margulis, *Probabilistic characteristics of graphs with large connectivity*, Problemy Peredači Informacii **10** (1974), 101–108.

[P1] I. Pak, *Random walks on groups: strong uniform time approach*, Ph.D. Thesis, Harvard U., 1997.

[P2] I. Pak, *When and how n choose k*, DIMACS Workshops on Randomized Algorithms (1998), AMS, Providence.

[P3] I. Pak, *Evolution of random walks on S_n,* in preparation (1998).

[P4] I. Pak, *Random walks on finite groups with few random generators*, Electr. J. Prob. **4** (1999), 1–11.

[R] Y. Roichman, *On random random walks*, Ann. Prob. **24** (1996), 1001–1011.

[S] R. P. Stanley, *Enumerative Combinatorics,* Vol. 1, Wadsworth & Brooks/Cole, California, 1986.

[T] M. Talagrand, *On Russo's approximative zero-one law*, Ann. Prob. **22** (1994), 1576–1587.

[WW] E. T. Whittaker, G. N. Watson, *A Course of Modern Analysis* (Fourth Edition), Cambridge University Press, Cambridge, UK, 1927.

[W] D. Wilson, *Random random walks on $\mathbb{Z}_2^n$*, Probab. Theory Related Fields **108** (1997), 441–457.