Expanders via Local Edge Flips in Quasilinear Time

George Giakkoupis
Inria, Rennes, France

May 25, 2022

Abstract

Mahlmann and Schindelhaue [24] proposed the following simple process, called *flip-chain*, for transforming any given connected d-regular graph into a d-regular expander: In each step, a random 3-path $abcd$ is selected, and edges ab and cd are replaced by two new edges ac and bd, provided that ac and bd do not exist already. A motivation for the study of the flip-chain arises in the design of overlay networks, where it is common practice that adjacent nodes periodically exchange random neighbors, to maintain good connectivity properties. It is known that the flip-chain converges to the uniform distribution over connected d-regular graphs, and it is conjectured that an expander graph is obtained after $O(nd \log n)$ steps, w.h.p., where n is the number of vertices. However, the best known upper bound on the number of steps is $O(n^2 d^2 \sqrt{\log n})$ [1], and the best bound on the mixing time of the chain is $O(n^{16} d^{36} \log n)$ [11, 6].

We provide a new analysis of a natural flip-chain instantiation, which shows that starting from any connected d-regular graph, for $d = \Omega(\log^2 n)$, an expander is obtained after $O(nd \log^2 n)$ steps, w.h.p. This result is tight within logarithmic factors, and almost matches the conjectured bound. Moreover, it justifies the use of edge flip operations in practice: for any d-regular graph with $d = \text{poly}(\log n)$, an expander is reached after each vertex participates in at most $\text{poly}(\log n)$ operations, w.h.p. Our analysis is arguably more elementary than previous approaches. It uses the novel notion of the *strain* of a cut, a value that depends both on the crossing edges and their adjacent edges. By keeping track of the cut strains, we form a recursive argument that bounds the time before all sets of a given size have large expansion, after all smaller sets have already attained large expansion.

1 Introduction

In [24], Mahlmann and Schindelhaue proposed a very simple and elegant process to transform any given connected d-regular graph into a d-regular expander.[1] This process consists of a sequence of *flip operations*. A flip operation on graph $G = (V, E)$ chooses a 3-path $abcd$ of G u.a.r., and if neither of the edges ac and bd exist already, then edges ab and cd are removed, and are replaced by edges ac and bd; otherwise, the operation does not modify the graph.

A flip operation does not change the degrees of vertices, and does not disconnect a connected graph. Moreover, it is a very *local* operation, as it affects only four vertices, at distance at most three apart. This is minimal, in the sense that no edge switching operation involving fewer than four vertices preserves the degrees, and the only shallower subgraph than a 3-path is the 3-star, for which there are no degree-preserving operations [6].

The Markov chain $\big(G_t = (V, E_t)\big)_{t \in \mathbb{N}_0}$ induced by a sequence of flip operations is called a *flip-chain*, and it converges to the uniform distribution over all connected d-regular graphs on V, if

[1] Both the input and output graphs, and also all graphs we consider throughout the paper, are simple, i.e., they have no loops or parallel edges. All graphs are also unweighted and undirected.

the initial graph $G_0 = (V, E_0)$ is connected and d-regular [24]. Moreover, based on experimental evidence, it has been conjectured that $t = O(nd \log n)$ operations suffice to ensure that graph G_t is an expander w.h.p. [24, 23].

A motivation for the study of flip-chains arises in the design of unstructured overlay networks (or peer-to-peer networks). Overlay networks should typically have low degree, small diameter and good connectivity, and should maintain these properties in the face of node arrivals and departures. Regular expanders, and in particular random regular graphs, possess all these desired properties [32]. Several unstructured overlay network designs have been proposed that build and maintain (near) regular expander graphs [30, 21, 12, 25, 13]. Some of these designs specifically aim at generating topologies that converge to random regular graphs [24, 5]. Random unstructured networks are also used in practical overlay libraries such as JXTA [29], where randomized edge swap operations between adjacent nodes are supported and commonly used to maintain a random topology [18, 11]. The analysis of flip-chains can thus provide some theoretical foundation for the heuristics used in practical unstructured overlays.

The flip-chain is a local variant of another well-known process, the *switch-chain*, proposed by McKay [26]. A switch operation chooses a pair of non-adjacent edges ab and cd u.a.r., and then, similarly to a flip operation, it replaces edges ab and cd by edges ac and bd, as long as ac and bd do not exist already. When started from a (not-necessarily connected) d-regular graph, the switch-chain converges to the uniform distribution over all d-regular graphs (on the same set of vertices). Thus, the switch-chain provides a natural process for generating a random d-regular graph. However, it is not suitable for the construction of overlay networks in a decentralized setting, because sampling a random pair of edges is a highly non-local operation, and, moreover, a switch operation can potentially disconnect a connected graph.

Both the switch-chain and the flip-chain have been studied extensively, for various families of graphs. Here, we focus on results for d-regular (undirected) graphs; for results on graphs with general degree sequences and directed graphs see [14, 10, 15, 2, 8, 9, 11, 27]. A restriction of the switch-chain on regular bipartite graphs was first analyzed in [19], using a canonical path argument. Very recently, in [31], a mixing time bound of $O((nd)^2 \log n)$ was shown for d-regular bipartite graphs, and a bound of $O(n \log^2 n)$ for the case of constant d, using a Markov chain comparison argument with the configuration model. For general d-regular graphs, a bound of $O(n^9 d^{24} \log n)$ was shown for the mixing time of the switch-chain [5, 4], using the canonical path method. Also, a sharp bound of $O(nd)$ was show in [1] for the time until an expander graph is obtained w.h.p. This result uses spectral techniques, in particular, it relies on the analysis of a potential function based on the Laplacian matrix exponential.

For the flip-chain on d-regular graphs, it was first shown that it is rapidly-mixing in [11], where an upper bound of $O(n^{53} d^{64} \log n)$ was derived for the mixing time, using a comparison argument with the switch-chain, and the result from [5]. Subsequently, an improved bound of $O(n^{16} d^{36} \log n)$ was shown in [7, 6], using a more refined canonical path argument. Also, [1] showed a bound of $O(n^2 d^2 \sqrt{\log n})$ for the time until and expander is obtained w.h.p., using spectral techniques.

All the above results, with the exception of the $O(n \log^2 n)$ bound shown in [31] and the $O(nd)$ bound shown in [1], are likely to be far from tight. For the flip-chain, in particular, even the best known bound of $O(n^2 d^2 \sqrt{\log n})$ until an expander is obtained, is almost quadratically larger than $O(nd \log n)$, which is generally believed to be the true bound.

1.1 Our Contribution

We show the first quasilinear[2] upper bound on the number of steps until an expander is obtained w.h.p., in a flip-chain starting from any connected d-regular graph, for $d = \Omega(\log^2 n)$. We consider the version of the flip operation analyzed in [1], which is slightly different than the original definition of [24], as explained below. The formal statement of the result is as follows.

Theorem 1. *Let $(G_t)_{t\geq 0}$ be a flip-chain, where G_0 is a connected d-regular graph with n vertices, and $d \geq 90\log^2 n$. Then, for any $t \geq 150\, nd\log^2 n$, the conductance of graph G_t satisfies that*

$$\Pr[\phi(G_t) \geq 1/32] = 1 - O(1/n^2).$$

This is a significant improvement over the previous best known bound of $O(n^2 d^2\sqrt{\log n})$ [1].[3] This is clear in the context of distributed overlay networks: Theorem 1 implies that an expander graph is obtained after each vertex participates in $O(d\log^2 n)$ flip operations, which is $\text{poly}(\log n)$ if $d = \text{poly}(\log n)$ (as is often the case in overlay networks). The previous bound, just ensures that the number of operations per vertex is at most $\tilde{O}(nd^2) = \text{poly}(n)$. Therefore, our result is the first to support that the flip-chain provides an *efficient* mechanism for distributed overlay construction of expander networks, and thus justifies the use of flip operations in practice.

The $O(nd\log^2 n)$ bound of Theorem 1 is nearly tight: In the ring-of-cliques graph [23], all the $\Theta(n/d)$ edges between cliques, except possibly for one, must be selected at least once by some flip operation, otherwise two unselected edges will form a cut-set of size 2 (see related Definition 13 and Lemma 14). And $\Omega(nd\log(n/d))$ many operation are needed for all those edges to be selected.

Note also that the bound is just by a factor $O(\log^2 n)$ larger than the $O(nd)$ bound for the switch-chain, which is not a local process. Moreover, it is comparable to time bounds for more structured distributed constructions of expanders, such as the skip graph construction in [16], which converts any connected n-vertex graph into one that contains a constant-degree expander as a subgraph [3], in $O(\log^2 n)$ parallel rounds. However, such structured constructions are significantly more involved.

As mentioned above, we assume a slightly different version of the flip process, used also in [1]. In the original flip operation, the selection process can be described as: first choose a random edge bc, then a random neighbor a of b, and a random neighbor d of c. In the variant we consider, edge bc and vertex a are chosen as before, but then, if a is not adjacent to c, vertex d is chosen at random among the neighbors of c *not adjacent to* b (if a is adjacent to c, the operation has no effect, as before). (See also Definition 11.) This modification may increase the probability that the operation succeeds to modify the graph, by at most a factor equal to the degree d (precisely, given bc, the success probability is the square-root of the original flip's). To implement the modified flip operation in a distributed system, we can either use multiple rounds of communication between vertices b and c, as proposed in [1], or have b send all its neighbors to c in a single round. It is easy to see that the first approach requires just a constant expected number of communication rounds. For the second approach, we remark that it is common in practice that adjacent nodes in an overlay network periodically send to each other their full list of neighbors, for maintenance purposes [18].

Our choice to adopt the above flip variant was driven by the computation of the expected strain change of a cut during a flip, in Lemma 17. (We define the *strain* of a cut below.) If one obtains a similar lemma for the original flip definition, then the rest of the analysis carries over with no changes.

The analysis of [1] requires that $d = \Omega(\log n)$, whereas we need that $d = \Omega(\log^2 n)$. As explained in Section 1.2, this is only necessary for the first part of the analysis, which bounds the number

[2] Quasilinear in the number of edges $m = nd/2$.

[3] However, the bound of [1] applies for $d = \Omega(\log n)$ rather than just for $d = \Omega(\log^2 n)$.

of steps until we obtain a graph with edge-connectivity $\Omega(d)$. This part of the analysis is also responsible for one of the two logarithmic factors in the time bound. Thus, if one could show that edge-connectivity $\Omega(d)$ is achieved in $O(nd \log n)$ steps for $d = \Omega(\log n)$, then our current analysis would yield the same result for the time until an expander is obtained. Obtaining similar bounds for the case of $d = O(1)$ is an important open problem.

Theorem 1 does not directly imply any bound on the mixing time of the flip-chain. It would be interesting to see whether it can be combined with existing analyses for the mixing time, to obtain better mixing-time bounds.

1.2 Techniques and Analysis Overview

Our analysis is different than previous analyses, which either use Markov chain comparison arguments and canonical path arguments [11, 6], or use techniques from linear algebra [1]. We use just elementary combinatorial/probabilistic tools.

A key element of our analysis is the novel definition of the *strain* of a cut. The strain of a *vertex* u with respect to a given cut is a value in the range $[0, 1/4]$, which is larger the more evenly the neighbors of u are distributed among the two sides of the cut.[4] The strain of the *cut* is then the sum of the strains over all vertices. Precisely, for a cut $\{S, \bar{S}\}$ in graph $G = (V, E)$, where $\bar{S} = V \setminus S$, the strain of the cut is $\sigma(S) = \sum_{u \in V} \alpha_u(1 - \alpha_u)$, where α_u is the fraction of neighbors of u in S. It is easy to see that if G is d-regular then $\sigma(S) \leq 2c(S)/d$, where $c(S) = |E(S, \bar{S})|$ is the cut size (Lemma 4). However, it is possible that $\sigma(S)$ is much smaller than $2c(S)/d$, e.g., it is zero if G is a bipartite graph with parts S and $\bar{S}$.

Before we describe how we use this new quantity, let us see what a most natural attempt to prove Theorem 1 would be. One would hope that it suffices to track the size of a given low-conductance cut $\{S, \bar{S}\}$, show that $c(S)$ increases in expectation in each step (at a sufficient rate), then show that $c(S)$ reaches $\Omega(d \cdot |S|)$ and remains high, with a probability sufficiently large that allows us to apply a union bound over all possible cuts. This simple plan fails for two reasons:

(1) The cut size $c(S)$ does not always increase in expectation at each step; it may decrease even if its conductance is small (as already pointed out in [1]).

(2) The total number of cuts is too large for a naively applied union bound to work.

Tracking the strain of a cut (instead of its size) will help us overcome issue (1). For issue (2), we will devise smarter ways to apply union bounds, over small enough sets of cuts.

Our proof consists of two parts, discussed next. The first part bounds the time until the min-cut size becomes at least $d/2$; the second part bounds the additional time until the graph conductance becomes at least $1/32$.

Edge-Connectivity Analysis. In this part, issue (1) above does not apply: We show that for any cut $\{S, \bar{S}\}$, if $c(S) \leq 3d/4$ (in general, at most $d - \epsilon$), then the expected increase of $c(S)$ in a step is sufficiently large (Lemma 16). This allows us to show that in $O(nd \log n)$ steps the cut size becomes $3d/4$ and remains at least $3d/4 - \lambda$, where $\lambda = \Theta(\log n)$, for $\mathrm{poly}(n)$ many rounds, w.h.p. (Lemmas 29 and 30).

Clearly, issue (2) *does* apply: we cannot just use a union bound over all the up to $n^{\Theta(d)}$ cuts of size at most $O(d)$. Instead, we show a basic structural lemma stating that if there are no sets S with $|S| \leq s$ and $c(S) \leq k$, then there are at most n sets T with $|T| \leq 2s$ and $c(T) \leq k$ (Lemma 8). This allows us to use a union bound over just $\mathrm{poly}(n)$ many sets (at most n new sets at each step),

[4] Intuitively, having neighbors on both sides of a cut is perceived as a source of (psychological) strain for the vertex.

to argue about all sets of size between s and $2s$, once we have dealt with sets of size smaller than s (Lemma 31).

Since we need to iterate the above argument logarithmically many times, the total time bound we obtain is $O(nd\log^2 n)$. Moreover, to be able to afford losing one term of $\lambda = \Theta(\log n)$ from $3d/4$ with each iteration, we need that $d = \Omega(\log^2 n)$. (See Lemma 28.)

Expansion Analysis. As soon as the size of a cut becomes at least d, issue (1) crops up. Then, rather than analyzing the cut size directly, we analyze its strain. The expected increase in a step of the strain $\sigma(S)$ of cut $\{S, \bar{S}\}$ consists of two terms: the first term is positive and is roughly [5]

$$(4/md) \cdot \sum_{\{u,v\} \in E} (a_u - a_v)^2;$$

the second term is negative but is negligible for the setting we will apply the result (key Lemma 17). We apply this result to graphs that are (s, φ)-*expanders*, for some constant φ, i.e., for every cut $\{S, \bar{S}\}$ of the graph, $c(S) \geq \varphi d \cdot \min\{|S|, |\bar{S}|, s\}$. (In particular, this implies that if $|S| \leq s$ then the conductance of the cut is at least φ.) For such graphs, we show that the sum above is $\Omega(c(S))$, thus the expected increase in the strain is $\Omega(c(S)/md)$ (see key Lemmas 23 and 27).

With this result in hand, we can analyze the evolution of the strain of a single cut, similarly to the size of a single cut in the Edge-Connectivity Analysis part. Roughly, as long as the graph remains an (s, φ)-expander, each the following holds with probability $1 - e^{-\Omega(sd)}$, for any given cut $\{S, \bar{S}\}$ with $s \leq |S| \leq n/2$: (i) the cut strain $\sigma(S)$ becomes at least $s/4$ in $O(m)$ steps; (ii) once that happens, the cut strain remains at least $s/8$ for poly(n) steps; and (iii) if the cut size is at least γsd, for a large constant γ, then the cut strain becomes at least $s/4$ before the cut size drops to sd (Lemmas 33 to 35). Recall that a lower bound on a cut strain immediately yields a lower bound for the cut size (and thus for the cut conductance), since always $\sigma(S) \leq 2c(S)/d$.

To deal with issue (2), we use a classic result by Karger [20], which bounds by $O(n^{2a})$ the number of cuts of size at most a times larger than the minimal (Theorem 10). We stress that the sole purpose of the preceding Edge-Connectivity Analysis part, which ensures that the min-cut size is $\Omega(d)$, is to allow us to use Karger's result with exponent $O(k/d)$, rather than $O(k)$, for cuts of size k. Indeed, the total number of cuts of size $O(sd)$ is then $n^{O(s)}$. Since the probability of our lemmas for a single cut is $1 - e^{-\Omega(sd)}$, we can afford to use a union bound over $n^{O(s)}$ cuts, as long as $d = \Omega(\log n)$. The reason we just need to consider cuts of size $c(S) = O(sd)$ is result (iii) we saw above. Without that, we would have to bound the number of all cuts of strain $\sigma(S) \leq s/4$, so we could not apply Karger's result directly (or even indirectly, as cuts of large size may still have a small strain).

We can then argue about multiple cuts, simultaneously, using union bounds (Lemma 36), and we finally obtain that, after achieving edge-connectivity $\Omega(d)$, the total number of steps to achieve constant conductance is $O(nd\log n)$, w.h.p. (Lemma 32).

Unlike the Edge-Connectivity Analysis part, the analysis of this part does not lose a logarithmic factor, and it only requires that $d = \Omega(\log n)$. However, the constant factors involved are large.

2 Definitions and Notation

For any graph $G = (V, E)$ and vertex $u \in V$, we denote by $\Gamma_G(u)$ the set of vertices adjacent to u in G, i.e., $\Gamma_G(u) = \{v \colon \{u, v\} \in E\}$. We also define $\Gamma^+_G(u) = \Gamma_G(u) \cup \{u\}$. If $\Gamma_G(u) \neq \emptyset$ (i.e., vertex

[5] Recall that a_u is the fraction of neighbors of u in set S. Also $m = nd/2$ is the number of edges.

u is not isolated), then for any set $S \subseteq V$ we denote by $\alpha_{G,u}(S)$ the fraction of u's neighbors that belong S, i.e.,

$$\alpha_{G,u}(S) = |S \cap \Gamma_G(u)|/|\Gamma_G(u)|.$$

For any two sets $S, T \subseteq V$, we let $E_G(S,T)$ be the set of edges $\{\{u,v\} : \{u,v\} \in E,\ u \in S, v \in T\}$.

We will often omit subscript G from the notations above, when there is not danger of confusion. Throughout the paper, n will denote the number of vertices of the graph considered, and m the number of edges. E.g., if the graph is $G = (V,E)$, then $n = |V|$ and $m = |E|$. If the graph is d-regular then $m = nd/2$, and we define $\delta = 1/d$.

2.1 Cut-Related Definitions

Recall that a cut of graph $G = (V,E)$ is a bipartition $\{S, \bar{S}\}$ of the set of vertices V, where both sets S and $\bar{S} = V \setminus S$ are non-empty.

Definition 2 (Cut Size)**.** The *size* of cut $\{S, \bar{S}\}$ in graph $G = (V,E)$ is the number of edges crossing the cut, and is denoted $c_G(S)$, i.e.,

$$c_G(S) = |E_G(S, \bar{S})|.$$

The *minimum* cut size over all cuts of G is denoted $c(G)$, and is called the *edge-connectivity* of G.

We introduce the related notion of the *strain* of a cut.

Definition 3 (Cut Strain)**.** The *strain* $\sigma_G(S)$ of cut $\{S, \bar{S}\}$ in graph $G = (V,E)$ is

$$\sigma_G(S) = \sum_{u \in V} \alpha_{G,u}(S) \cdot \alpha_{G,u}(\bar{S}).$$

Note that $c_G(S) = c_G(\bar{S})$ and $\sigma_G(S) = \sigma_G(\bar{S})$. The following simple relation holds between the two quantities.

Lemma 4. *If $G = (V,E)$ is a d-regular graph then, for any cut $\{S, \bar{S}\}$, $\sigma_G(S) \le 2c_G(S)/d$.*

Proof. We have

$$\sigma_G(S) \le \sum_{u \in S} \alpha_u(\bar{S}) + \sum_{u \in \bar{S}} \alpha_u(S) = 2c_G(S)/d.$$

□

Definition 5 (Conductance)**.** The *conductance of cut* $\{S, \bar{S}\}$ in d-regular graph $G = (V,E)$ is

$$\phi_G(S) = c_G(S) / \left(d \cdot \min\{|S|, |\bar{S}|\}\right).$$

The *conductance of graph* G is $\phi(G) = \min\{\phi_G(S) : \emptyset \subset S \subset V\}$.

We will need also the following definitions for the analysis.

Definition 6 ($c(G,s)$ and $\mathcal{C}_G(s,k)$)**.** For any graph $G = (V,E)$,

$$c(G,s) = \min\{c_G(S) : S \subseteq V,\ 1 \le |S| \le s\}$$
$$\mathcal{C}_G(s,k) = \{S : S \subseteq V,\ 1 \le |S| \le s,\ c_G(S) \le k\}.$$

Note that $c(G,s) = \min\{k : \mathcal{C}_G(s,k) \ne \emptyset\}$, and $c(G,s) = c(G)$ if $s \ge \lfloor n/2 \rfloor$.

Definition 7 (((s,φ)-Expander)**.** A d-regular graph $G=(V,E)$ is an (s,φ)-*expander* if for every cut $\{S,\bar{S}\}$, where $|S|\le|\bar{S}|$,

$$c_G(S)\ge\varphi d\cdot\min\{|S|,s\}.$$

The set of all (s,φ)-expander graphs is denoted $\mathbb{G}_{s,\varphi}$.

Our definition of an (s,φ)-expander is stronger than the common definition of an (s,φ)-small-set expander, which only requires that $c_G(S)\ge\varphi d\cdot|S|$ if $|S|\le s$, and allows for the cut size to be arbitrary small if $s<|S|\le n/2$.

The following simple facts hold for any d-regular graph G. If G is an (s,φ)-expander, then $c(G)\ge\varphi d$. If $c(G)\ge k$, then G is a $(1,k/d)$-expander. If G is an (s,φ)-expander and $s\ge\lfloor n/2\rfloor$, then the graph conductance is $\phi(G)\ge\varphi$.

3 Upper Bounds on the Number of Cuts

In this section, we provide upper bounds on the number of cuts whose size is at most equal to a given value.

The next structural lemma bounds the cardinality of the set $\mathcal{C}_G(2s+1,k)$ when $c(G,s)>k$ (or, equivalently, when $\mathcal{C}_G(s,k)=\emptyset$).

Lemma 8. *For any regular graph $G=(V,E)$, if $c(G,s)>k$ then $|\mathcal{C}_G(2s+1,k)|\le n$.*

Proof. Let $s'=2s+1$. Let M denote the set containing all *minimal* sets $S\in\mathcal{C}_G(s',k)$. (S is minimal if it is not a proper subset of another set $S'\in\mathcal{C}_G(s',k)$.)

Claim 8.1. If $S_1,S_2\in M$ and $S_1\neq S_2$, then $S_1\cap S_2=\emptyset$.

Proof of Claim 8.1. Suppose, for contradiction that $S_1\cap S_2\neq\emptyset$. Then, since $S_1,S_2\in M$ and $S_1\neq S_2$, it follows that $S_1\backslash S_2\neq\emptyset$ and $S_2\backslash S_1\neq\emptyset$. Let $k_1=|E(S_1,S_2\backslash S_1)|$ and $k_2=|E(S_2,S_1\backslash S_2)|$, and w.l.o.g. suppose that $k_1\le k_2$. Then $c_G(S_2\setminus S_1)=c_G(S_2)-k_2+k_1\le c_G(S_2)\le k$, thus $S_2\setminus S_1\in\mathcal{C}_G(s',k)$. We thus conclude that set S_2 is not minimal – a contradiction. □

Claim 8.2. If $S_1,S_2,T\in\mathcal{C}_G(s',k)$, and $T\subseteq S_1\cap S_2$, then $S_1\subseteq S_2$ or $S_2\subseteq S_1$.

Proof of Claim 8.2. Suppose, for contradiction, that $S_1\nsubseteq S_2$ and $S_2\nsubseteq S_1$. Then $S_1\setminus S_2\neq\emptyset$ and $S_2\setminus S_1\neq\emptyset$. As in the proof of Claim 8.1, we let $k_1=|E(S_1,S_2\setminus S_1)|$ and $k_2=|E(S_2,S_1\setminus S_2)|$, and assume $k_1\le k_2$ to obtain $S_2\setminus S_1\in\mathcal{C}_G(s',k)$. From that and lemma's assumption $c(G,s)>k$, it follows $|S_2\backslash S_1|\ge s+1$. Also, $|S_1\cap S_2|\ge|T|\ge s+1$. Hence, $|S_2|=|S_2\backslash S_1|+|S_1\cap S_2|\ge 2(s+1)>s'$, which contradicts that $S_2\in\mathcal{C}_G(s',k)$. □

Claim 8.3. If $S\in\mathcal{C}_G(s',k)\setminus M$ then there is a set $T\in M$ such that $T\subset S$.

Proof of Claim 8.3. By definition, a non-minimal set has a *proper* subset in $\mathcal{C}_G(s',k)$. By iterating this argument at most s times, we obtain that S has a proper *minimal* subset. □

From Claim 8.1 and assumption $c(G,s)>k$, we have that the number of minimal sets is $|M|\le n/(s+1)$. From Claim 8.2, it follows that for each minimal set $T\in M$, the total number of non-minimal sets $S\in\mathcal{C}_G(s',k)\setminus M$ that contain T as a subset is at most $s'-|T|\le s'-(s+1)\le s$. Since, from Claim 8.3, each non-minimal set must contain a minimal set, we conclude that

$$|\mathcal{C}_G(s',k)|\le|M|\cdot(1+s)\le\frac{n}{s+1}\cdot(1+s)=n.$$ □

The condition of Lemma 8, that $c(G, s) > k$, is always true for $s = d$ and $k = d - 1$:

Lemma 9. *For any d-regular graph G, $c(G, d) \geq d$.*

Proof. The statement follows from the fact that if $|S| \leq d$ then each $u \in S$ has at least $d - |S| + 1$ neighbors in $\bar{S}$. □

We will also need the following standard result. An a-minimal cut is a cut of size within a multiplicative factor of a of the minimum.

Theorem 10 ([20])**.** *In any graph, the number of a-minimal cuts is $O\left(n^{\lfloor 2a \rfloor}\right)$, for any real $a \geq 1$.*

4 Flip Operations

We define a flip operation as follows. (This definition is equivalent to that of a *random-flip transformation* from [1].)

Definition 11 (Flip Operation)**.** A *flip* operation on graph $G = (V, E)$ consists of the following steps:

1. Choose an (ordered) pair of adjacent vertices $a, b \in V$ u.a.r.
2. Choose a vertex $a' \in \Gamma_G(a)$ u.a.r. (possibly, $a' = b$).
3. If the following two conditions hold:[6]

$$a' \in \Gamma_G(a) \setminus \Gamma_G^+(b) \quad \text{and} \quad \Gamma_G(b) \setminus \Gamma_G^+(a) \neq \emptyset, \tag{1}$$

   3.1. Choose a vertex $b' \in \Gamma_G(b) \setminus \Gamma_G^+(a)$ u.a.r.

   3.2. Output graph $\left(V, (E \setminus \{e_1, e_2\}) \cup \{e_1', e_2'\}\right)$, where $e_1 = \{a, a'\}$, $e_2 = \{b, b'\}$, $e_1' = \{a, b'\}$, and $e_2' = \{b, a'\}$, i.e., replace edges e_1 and e_2 of G with the new edges e_1' and e_2'.

4. Else (i.e., if (1) does not hold), output G.

We denote by $\text{Flip}(G)$ the output graph. We say the flip operation is *good* if (1) holds.

Flip operations do not change the degree of vertices, and do not disconnect a connected graph (more generally, they do not change the set of connected components of the graph).

The first step of a flip operation is equivalent to sampling an edge from E u.a.r., and then choosing a random order of its endpoints. In general, the order chosen *does* affect the distribution of the outcome, but in the case of regular graphs it does not, as shown in the next lemma.

Lemma 12 (Symmetry Lemma)**.** *If $G = (V, E)$ is a regular graph, then for any edge $\{u, v\} \in E$ and vertices $u' \in \Gamma_G(u) \setminus \Gamma_G^+(v)$ and $v' \in \Gamma_G(v) \setminus \Gamma_G^+(u)$,*

$$\Pr[(a, b, a', b') = (u, v, u', v')] = \Pr[(a, b, a', b') = (v, u, v', u')].$$

Proof. Let d be the degree of G. Then $\Pr[(a, b, a', b') = (u, v, u', v')]$ equals

$$\frac{1}{nd} \cdot \frac{1}{d} \cdot \frac{1}{|\Gamma_G(v) \setminus \Gamma_G^+(u)|} = \frac{1}{nd} \cdot \frac{1}{d} \cdot \frac{1}{|\Gamma_G(u) \setminus \Gamma_G^+(v)|},$$

which is equal to $\Pr[(a, b, a', b') = (v, u, v', u')]$. □

[6] If G is regular, then the first condition in (1) implies the second, because $|\Gamma_G(a) \setminus \Gamma_G^+(b)| = |\Gamma_G(b) \setminus \Gamma_G^+(a)|$.

A flip operation is *pertinent* to cut $\{S, \bar{S}\}$ if at least one of the edges $\{a, b\}$, $\{a, a'\}$, and $\{b, b'\}$ crosses the cut. The formal definition is as follows.

Definition 13 (Pertinent Flip)**.** A flip operation on $G = (V, E)$ is *pertinent to cut* $\{S, \bar{S}\}$, if $\{a, b\} \in E(S, \bar{S})$, or $\{a, a'\} \in E(S, \bar{S})$, or the flip is good and $\{b, b'\} \in E(S, \bar{S})$.

If a flip is not pertinent to a given cut, then the operation does not change the set of crossing edges of the cut, thus it does not change the cut's size or strain.

Lemma 14. *If $G = (V, E)$ is a regular graph then, for any cut $\{S, \bar{S}\}$, the probability of the event $\mathcal{P}_S$ that a flip is pertinent to the cut is*

$$\Pr[\mathcal{P}_S] \leq 3c_G(S)/m.$$

Proof. From the definition of a flip, it is immediate that $\Pr[\{a, b\} \in E(S, \bar{S})] = c_G(S)/m$ and $\Pr[\{a, a'\} \in E(S, \bar{S})] = c_G(S)/m$. Also, if $\mathcal{E}$ is the event that the flip is good then, from Lemma 12,

$$\Pr\left[\mathcal{E} \cap \left\{\{b, b'\} \in E(S, \bar{S})\right\}\right] = \Pr\left[\mathcal{E} \cap \left\{\{a, a'\} \in E(S, \bar{S})\right\}\right] \leq \Pr[\{a, a'\} \in E(S, \bar{S})] = c_G(S)/m.$$

We complete the proof by combining the above three equations using a union bound. □

A flip-chain is a Markov chain over graphs, with transitions corresponding to a single flip operation.

Definition 15 (Flip-Chain)**.** A *flip-chain* is a Markov chain $\left(G_t = (V, E_t)\right)_{t \geq 0}$, where $G_t = \mathrm{Flip}(G_{t-1})$, for each $t \geq 1$. We call the chain an (n, d)*-flip-chain* if $|V| = n$ and G_0 is a connected d-regular graph (thus all graphs G_t are connected and d-regular).

We only consider (n, d)-flip-chains in this paper. In our analysis of flip-chains, in Sections 6 and 7, we replace subscripts G_t in all notations by t, e.g., we write $c_t(S)$ and $\sigma_t(S)$ instead of $c_{G_t}(S)$ and $\sigma_{G_t}(S)$, respectively. For $t \geq 1$, we denote by flip_t the flip operation on G_{t-1} that yields G_t, and by a_t, b_t, a'_t, b'_t the vertices selected in that operation (where b'_t is defined only if flip_t is a good operation).

5 Expected Cut Size and Strain After a Single Flip

5.1 Additional Notation

For a connected graph $G = (V, E)$, recall that $\alpha_{G,u}(S) = |S \cap \Gamma_G(u)|/|\Gamma_G(u)|$, and for any edge $\{u, v\} \in E$, let

$$\alpha_{G,u,v}(S) = |S \cap \Gamma(u) \cap \Gamma(v)|/|\Gamma(u)|$$
$$\beta_{G,u,v}(S) = |S \cap \Gamma(u) \setminus \Gamma^+(v)|/|\Gamma(u)|$$
$$\rho_{G,u,v} = |\Gamma(u) \setminus \Gamma^+(v)|/|\Gamma(u)|.$$

In this section, we consider just a pair of graphs G, G', where G is given and $G' = \mathrm{Flip}(G)$. The set S is also fixed and given. We thus use the following shortcut notation throughout Section 5,

$$\alpha_u = \alpha_{G,u}(S),\ \bar{\alpha}_u = \alpha_{G,u}(\bar{S}),\ \alpha_{u,v} = \alpha_{G,u,v}(S),\ \bar{\alpha}_{u,v} = \alpha_{G,u,v}(\bar{S}),$$
$$\beta_{u,v} = \beta_{G,u,v}(S),\ \bar{\beta}_{u,v} = \beta_{G,u,v}(\bar{S}),\ \rho_{u,v} = \rho_{G,v,u}.$$

Also, $\Gamma(u) = \Gamma_G(u)$ and $\Gamma^+(u) = \Gamma^+_G(u)$. We use primed notation to denote the same quantities in G', e.g., $\bar{\alpha}'_u = \alpha_{G',u}(\bar{S})$ and $\Gamma'(u) = \Gamma_{G'}(u)$.

We will frequently use the following simple equations,[7]

$$\alpha_u + \bar{\alpha}_u = 1, \quad \beta_{u,v} = \alpha_u - \alpha_{u,v} - \mathbb{1}_{v \in S}/|\Gamma(u)|, \quad \bar{\beta}_{u,v} = \bar{\alpha}_u - \bar{\alpha}_{u,v} - \mathbb{1}_{v \in \bar{S}}/|\Gamma(u)|$$
$$\rho_{u,v} = 1 - \alpha_{u,v} - \bar{\alpha}_{u,v} - 1/|\Gamma(u)| = \beta_{u,v} + \bar{\beta}_{u,v}.$$

Since G will be a d-regular graph, we also have that $\alpha_{u,v} = \alpha_{v,u}$, $\bar{\alpha}_{u,v} = \bar{\alpha}_{v,u}$, $\rho_{u,v} = \rho_{v,u}$, and $\rho_{u,v} = 0$ if and only if $\Gamma^+(u) = \Gamma^+(v)$.

Here are some examples that use the above notation. Suppose that $(a, b) = (u, v)$ in a flip operation. Then, the probability the flip is good is $\rho_{u,v} = \rho_{v,u}$. The probability the flip is good and $a' \in S$ is $\beta_{u,v}$. The probability the flip is good and $a' \in \bar{S}$ is $\bar{\beta}_{u,v}$. The probability the flip is good and $a', b' \in S$ is $\beta_{u,v} \cdot (\beta_{v,u}/\rho_{v,u}) = \beta_{u,v}\beta_{v,u}/\rho_{u,v}$.

5.2 Expected Size of a Small Cut After a Flip

We give a simple lower bound on the expected increase in the size of a given cut $\{S, \bar{S}\}$, after a single flip operation is applied to d-regular graph G. The bound is useful only if the initial size of the cut is small, namely, $c_G(S) \leq d$. Note that the cut size does not change unless $\{a, b\} \in E_G(S, \bar{S})$, thus we condition the expectation on the event $\{a, b\} = \{u, v\} \in E_G(S, \bar{S})$. Recall that $\delta = 1/d$.

Lemma 16. *Let $G = (V, E)$ be a d-regular graph, let $\{S, \bar{S}\}$ be any cut, and let $\{u, v\} \in E_G(S, \bar{S})$. If $G' = \mathrm{Flip}(G)$ then*

$$\mathbf{E}[c_{G'}(S) \mid \{a, b\} = \{u, v\}] \geq c_G(S) + 2(1 - \delta c_G(S)).$$

Proof. Suppose w.l.o.g. that $u \in S$ and $v \in \bar{S}$. If $(a, b) = (u, v)$ then: $c_{G'}(S) - c_G(S) = 2$ when the flip is good and $a' \in S$, $b' \in \bar{S}$, which happen with probability $\beta_{u,v}\bar{\beta}_{v,u}/\rho_{u,v}$; $c_{G'}(S) - c_G(S) = -2$ when the flip is good and $a' \in \bar{S}$, $b' \in S$, which happen with probability $\bar{\beta}_{u,v}\beta_{v,u}/\rho_{u,v}$; and $c_{G'}(S) - c_G(S) = 0$ in all other cases. Then,

$$\begin{aligned}
\mathbf{E}[c_{G'}(S) \mid (a, b) = (u, v)] - c_G(S) &= 2\beta_{u,v}\bar{\beta}_{v,u}/\rho_{u,v} - 2\bar{\beta}_{u,v}\beta_{v,u}/\rho_{u,v} \\
&= 2\beta_{u,v}(\rho_{u,v} - \beta_{v,u})/\rho_{u,v} - 2(\rho_{u,v} - \beta_{u,v})\beta_{v,u}/\rho_{u,v} \\
&= 2(\beta_{u,v} - \beta_{v,u}) \\
&= 2(\alpha_u - \alpha_v + \delta) \\
&= 2(1 - (\bar{\alpha}_u + \alpha_v - \delta)) \\
&\geq 2(1 - \delta c_G(S)).
\end{aligned}$$

Also, from Lemma 12, it follows $\mathbf{E}[c_{G'}(S) \mid (a, b) = (v, u)] = \mathbf{E}[c_{G'}(S) \mid (a, b) = (u, v)]$. ◻

5.3 Expected Cut Strain After a Flip

In this section, we compute the expected change in the strain of a given cut $\{S, \bar{S}\}$, after a single flip operation is applied to d-regular graph G. Recall $\delta = 1/d$.

Lemma 17. *If $G = (V, E)$ is a d-regular graph and $G' = \mathrm{Flip}(G)$, then for any cut $\{S, \bar{S}\}$,*

$$\mathbf{E}[\sigma_{G'}(S)] = \sigma_G(S) + 4\delta m^{-1} \sum_{\{u,v\} \in E} (\beta_{u,v} - \beta_{v,u})^2 - \delta^2 m^{-1} \sum_{\{u,v\} \in E} \gamma_{G,\{u,v\}}(S),$$

[7] By $\mathbb{1}_{\mathcal{E}}$ we denote the indicator function that is 1 if $\mathcal{E}$ holds and is 0 otherwise.

where $\gamma_{G,\{u,v\}}(S) = 0$ *if* $\Gamma^+(u) = \Gamma^+(v)$, *otherwise*

$$\gamma_{G,\{u,v\}}(S) = \begin{cases} \bar{\beta}_{u,v} + \bar{\beta}_{v,u} + 2(\beta_{u,v}\bar{\beta}_{v,u} + \bar{\beta}_{u,v}\beta_{v,u})/\rho_{u,v} & \text{if } u, v \in S \\ \beta_{u,v} + \beta_{v,u} + 2(\beta_{u,v}\bar{\beta}_{v,u} + \bar{\beta}_{u,v}\beta_{v,u})/\rho_{u,v} & \text{if } u, v \in \bar{S} \\ \beta_{u,v} + \bar{\beta}_{v,u} + 4\beta_{u,v}\bar{\beta}_{v,u}/\rho_{u,v} & \text{if } u \in S, v \in \bar{S}. \end{cases}$$

For each vertex $u \in V$, let

$$y_u = \alpha'_u\bar{\alpha}'_u - \alpha_u\bar{\alpha}_u.$$

Then, $\sigma_{G'}(S) - \sigma_G(S) = \sum_{u \in V} y_u$. If the flip operation is good and $u \in \{a, b, a', b'\}$, then the sets $\Gamma(u)$ and $\Gamma'(u)$ differ in exactly one element; otherwise $\Gamma(u) = \Gamma'(u)$. Hence, if $\alpha'_u < \alpha_u$ then $\alpha'_u = \alpha_u - \delta$ and $\bar{\alpha}'_u = \bar{\alpha}_u + \delta$; and if $\alpha'_u > \alpha_u$ then $\alpha'_u = \alpha_u + \delta$ and $\bar{\alpha}'_u = \bar{\alpha}_u - \delta$. It is then easy to compute

$$y_u = \begin{cases} \delta(\alpha_u - \bar{\alpha}_u) - \delta^2 & \text{if } \alpha'_u < \alpha_u \\ \delta(\bar{\alpha}_u - \alpha_u) - \delta^2 & \text{if } \alpha'_u > \alpha_u \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

In the following subsections, we compute the expected values of y_a, y_b, $y_{a'}$, and $y_{b'}$, and then combine these results to compute the expected value of $\sigma_{G'}(S) - \sigma_G(S)$.

It is convenient to divide $\gamma_{G,\{u,v\}}(S)$ into two parts. For any $\{u, v\} \in E$, let

$$\eta_{\{u,v\}} = \begin{cases} \bar{\beta}_{u,v} + \bar{\beta}_{v,u} & \text{if } u, v \in S \\ \beta_{u,v} + \beta_{v,u} & \text{if } u, v \in \bar{S} \\ \beta_{u,v} + \bar{\beta}_{v,u} & \text{if } u \in S, v \in \bar{S} \end{cases}, \quad \zeta_{\{u,v\}} = \begin{cases} 2(\beta_{u,v}\bar{\beta}_{v,u} + \bar{\beta}_{u,v}\beta_{v,u})/\rho_{u,v} & \text{if } u, v \in S \text{ or } \bar{S} \\ 4\beta_{u,v}\bar{\beta}_{v,u}/\rho_{u,v} & \text{if } u \in S, v \in \bar{S}. \end{cases}$$

Then $\gamma_{G,\{u,v\}}(S) = \eta_{\{u,v\}} + \zeta_{\{u,v\}}$.

5.3.1 Expected Value of $y_a + y_b$

First, we consider the case where $\{a, b\} \in E_G(S, \bar{S})$.

Lemma 18. *For any edge* $\{u, v\} \in E$, *where* $u \in S$, $v \in \bar{S}$, *and* $\Gamma^+(u) \neq \Gamma^+(v)$,

$$\mathbf{E}[y_a + y_b \mid \{a, b\} = \{u, v\}] = 2\delta(\beta_{u,v} - \beta_{v,u})^2 - \delta^2\zeta_{\{u,v\}}.$$

Proof. When $(a, b) = (u, v)$, the values of y_a and y_b are non-zero only if (i) the flip is good and $a' \in S, b' \in \bar{S}$, or (ii) the flip is good and $a' \in \bar{S}, b' \in S$. By (2), if (i) holds then $y_a = \delta(2\alpha_u - 1) - \delta^2$ and $y_b = \delta(1 - 2\alpha_v) - \delta^2$; while if (ii) holds then $y_a = \delta(1 - 2\alpha_u) - \delta^2$ and $y_b = \delta(2\alpha_v - 1) - \delta^2$. Also, given $(a, b) = (u, v)$, the probability that (i) holds is $\beta_{u,v}\bar{\beta}_{v,u}/\rho_{u,v}$; while the probability that (ii) holds is $\bar{\beta}_{u,v}\beta_{v,u}/\rho_{u,v}$. Then,

$$\begin{aligned} \mathbf{E}[y_a + y_b \mid (a, b) = (u, v)] &= [\delta(2\alpha_u - 1) - \delta^2 + \delta(1 - 2\alpha_v) - \delta^2] \cdot \beta_{u,v}\bar{\beta}_{v,u}/\rho_{u,v} \\ &\quad + [\delta(1 - 2\alpha_u) - \delta^2 + \delta(2\alpha_v - 1) - \delta^2] \cdot \bar{\beta}_{u,v}\beta_{v,u}/\rho_{u,v} \\ &= [2\delta(\beta_{u,v} - \beta_{v,u}) - 4\delta^2] \cdot \beta_{u,v}\bar{\beta}_{v,u}/\rho_{u,v} \\ &\quad + [2\delta(\beta_{v,u} - \beta_{u,v})] \cdot \bar{\beta}_{u,v}\beta_{v,u}/\rho_{u,v} \\ &= 2\delta(\beta_{u,v} - \beta_{v,u})^2 - 4\delta^2\beta_{u,v}\bar{\beta}_{v,u}/\rho_{u,v}, \end{aligned}$$

where for the last equation we used $\rho_{u,v} = \beta_{u,v} + \bar{\beta}_{u,v} = \beta_{v,u} + \bar{\beta}_{v,u}$. Also, from Lemma 12, it follows $\mathbf{E}[y_a + y_b \mid (a, b) = (v, u)] = \mathbf{E}[y_a + y_b \mid (a, b) = (u, v)]$. ☐

Next, we prove a similar result for the case where $\{a,b\} \in E \setminus E_G(S,\bar{S})$. (Recall that the definition of $\zeta_{\{u,v\}}$ is different in this case.)

Lemma 19. *For any edge $\{u,v\} \in E$, where either $u,v \in S$ or $u,v \in \bar{S}$, and $\Gamma^+(u) \neq \Gamma^+(v)$,*

$$\mathbf{E}[y_a + y_b \mid \{a,b\} = \{u,v\}] = 2\delta(\beta_{u,v} - \beta_{v,u})^2 - \delta^2 \zeta_{\{u,v\}}.$$

Proof. The proof is similar to that of Lemma 18. By the same reasoning, we have

$$\begin{aligned}
\mathbf{E}[y_a + y_b \mid (a,b) = (u,v)] &= [\delta(2\alpha_u - 1) - \delta^2 + \delta(1 - 2\alpha_v) - \delta^2] \cdot \beta_{u,v}\bar{\beta}_{v,u}/\rho_{u,v} \\
&\quad + [\delta(1 - 2\alpha_u) - \delta^2 + \delta(2\alpha_v - 1) - \delta^2] \cdot \bar{\beta}_{u,v}\beta_{v,u}/\rho_{u,v} \\
&= [2\delta(\beta_{u,v} - \beta_{v,u}) - 2\delta^2] \cdot \beta_{u,v}\bar{\beta}_{v,u}/\rho_{u,v} \\
&\quad + [2\delta(\beta_{v,u} - \beta_{u,v}) - 2\delta^2] \cdot \bar{\beta}_{u,v}\beta_{v,u}/\rho_{u,v} \\
&= 2\delta(\beta_{u,v} - \beta_{v,u})^2 - 2\delta^2(\beta_{u,v}\bar{\beta}_{v,u} + \bar{\beta}_{u,v}\beta_{v,u})/\rho_{u,v}.
\end{aligned}$$

Also, $\mathbf{E}[y_a + y_b \mid (a,b) = (v,u)] = \mathbf{E}[y_a + y_b \mid (a,b) = (u,v)]$, since the formula in the last line is symmetric for u and v. □

5.3.2 Expected Values of $y_{a'}$ and $y_{b'}$

We first compute the expected value of $y_{a'}$ conditionally on $a' = u$, and then remove the condition.

Lemma 20. *For any vertex $u \in V$,*

$$\mathbf{E}[y_{a'} \mid a' = u] = \delta^2 \sum_{v \in \Gamma(u)} Y_{u,v} - \delta^3 \sum_{v \in \Gamma(u) \cap S} \bar{\beta}_{v,u} - \delta^3 \sum_{v \in \Gamma(u) \cap \bar{S}} \beta_{v,u},$$

where

$$Y_{u,v} = \begin{cases} (\alpha_u - \bar{\alpha}_u)(\bar{\alpha}_v - \delta \cdot \mathbb{1}_{u \in \bar{S}}) & \text{if } v \in S \\ (\bar{\alpha}_u - \alpha_u)(\alpha_v - \delta \cdot \mathbb{1}_{u \in S}) & \text{if } v \in \bar{S}. \end{cases}$$

Proof. For any vertex $v \in \Gamma(u) \cap S$, we have

$$\mathbf{E}[y_{a'} \mid (a',a) = (u,v)] = [\delta(\alpha_u - \bar{\alpha}_u) - \delta^2] \cdot \bar{\beta}_{v,u},$$

because $y_{a'} = \delta(\alpha_u - \bar{\alpha}_u) - \delta^2$ if the flip is good and $b \in \bar{S}$ as $\alpha'_u < \alpha_u$ in this case, and $y_{a'} = 0$ otherwise; and, given $(a',a) = (u,v)$, the event that the flip is good and $b \in \bar{S}$ is the same as the event that $b \in \bar{S} \setminus \Gamma^+(u)$, whose probability is $\bar{\beta}_{v,u}$. Similarly, we obtain that for any $v \in \Gamma(u) \cap \bar{S}$,

$$\mathbf{E}[y_{a'} \mid (a',a) = (u,v)] = [\delta(\bar{\alpha}_u - \alpha_u) - \delta^2] \cdot \beta_{v,u}.$$

It follows that

$$\begin{aligned}
\mathbf{E}[y_{a'} \mid a' = u] &= \sum_{v \in \Gamma(u) \cap S} [\delta^2(\alpha_u - \bar{\alpha}_u) - \delta^3] \cdot \bar{\beta}_{v,u} \\
&\quad + \sum_{v \in \Gamma(u) \cap \bar{S}} [\delta^2(\bar{\alpha}_u - \alpha_u) - \delta^3] \cdot \beta_{v,u} \\
&= \delta^2 \sum_{v \in \Gamma(u) \cap S} (\alpha_u - \bar{\alpha}_u)(\bar{\alpha}_v - \bar{\alpha}_{u,v} - \delta \cdot \mathbb{1}_{u \in \bar{S}}) - \delta^3 \sum_{v \in \Gamma(u) \cap S} \bar{\beta}_{v,u} \\
&\quad + \delta^2 \sum_{v \in \Gamma(u) \cap \bar{S}} (\bar{\alpha}_u - \alpha_u)(\alpha_v - \alpha_{u,v} - \delta \cdot \mathbb{1}_{u \in S}) - \delta^3 \sum_{v \in \Gamma(u) \cap \bar{S}} \beta_{v,u}.
\end{aligned}$$

Finally, the terms $\bar{\alpha}_{u,v}$ and $\alpha_{u,v}$ in the last two lines get cancelled, as it is easy to check that

$$\sum_{v\in\Gamma(u)\cap S}\bar{\alpha}_{u,v}=\sum_{v\in\Gamma(u)\cap\bar{S}}\alpha_{u,v}.$$

□

Lemma 21. $\mathbf{E}[y_{a'}]=\delta m^{-1}\sum_{\{u,v\}\in E}(\beta_{u,v}-\beta_{v,u})^2-\delta^2 m^{-1}\sum_{\{u,v\}\in E}\eta_{\{u,v\}}/2.$

Proof. From Lemma 20,

$$\mathbf{E}[y_{a'}]=n^{-1}\sum_{u\in V}\left(\delta^2\sum_{v\in\Gamma(u)}Y_{u,v}-\delta^3\sum_{v\in\Gamma(u)\cap S}\bar{\beta}_{v,u}-\delta^3\sum_{v\in\Gamma(u)\cap\bar{S}}\beta_{v,u}\right).$$

We have

$$\sum_{u\in V}\sum_{v\in\Gamma(u)}Y_{u,v}=\sum_{\{u,v\}\in E}(Y_{u,v}+Y_{v,u}).$$

We now compute the sum $Y_{u,v}+Y_{v,u}$, for each edge $\{u,v\}\in E$. We have the following three cases: If $u,v\in S$,

$$\begin{aligned}Y_{u,v}+Y_{v,u}&=(\alpha_u-\bar{\alpha}_u)\bar{\alpha}_v+(\alpha_v-\bar{\alpha}_v)\bar{\alpha}_u\\&=(\alpha_u-\bar{\alpha}_u)(\bar{\alpha}_u+(\bar{\alpha}_v-\bar{\alpha}_u))+(\alpha_v-\bar{\alpha}_v)(\bar{\alpha}_v+(\bar{\alpha}_u-\bar{\alpha}_v))\\&=(\alpha_u-\bar{\alpha}_u)\bar{\alpha}_u+(\alpha_v-\bar{\alpha}_v)\bar{\alpha}_v+2(\beta_{u,v}-\beta_{v,u})^2.\end{aligned}$$

If $u,v\in\bar{S}$, we obtain similarly

$$Y_{u,v}+Y_{v,u}=(\bar{\alpha}_u-\alpha_u)\alpha_u+(\bar{\alpha}_v-\alpha_v)\alpha_v+2(\beta_{u,v}-\beta_{v,u})^2.$$

Last, if $u\in S$ and $v\in\bar{S}$, we get

$$\begin{aligned}Y_{u,v}+Y_{v,u}&=(\bar{\alpha}_u-\alpha_u)(\alpha_v-\delta)+(\alpha_v-\bar{\alpha}_v)(\bar{\alpha}_u-\delta)\\&=(\bar{\alpha}_u-\alpha_u)\alpha_u+(\alpha_v-\bar{\alpha}_v)\bar{\alpha}_v+2(\beta_{u,v}-\beta_{v,u})^2-2\delta(\beta_{u,v}-\beta_{v,u}).\end{aligned}$$

Summing over all $\{u,v\}\in E$, results in the cancellation of all terms $(\alpha_x-\bar{\alpha}_x)\bar{\alpha}_x$ and $(\bar{\alpha}_x-\alpha_x)\alpha_x$, for $x\in\{u,v\}$, and we obtain

$$\sum_{\{u,v\}\in E}(Y_{u,v}+Y_{v,u})=\sum_{\{u,v\}\in E}2(\beta_{u,v}-\beta_{v,u})^2-\sum_{u\in S}\sum_{v\in\Gamma(v)\cap\bar{S}}2\delta(\beta_{u,v}-\beta_{v,u}).$$

It follows that

$$\begin{aligned}\mathbf{E}[y_{a'}]=n^{-1}\Bigg(&2\delta^2\sum_{\{u,v\}\in E}(\beta_{u,v}-\beta_{v,u})^2-\delta^3\sum_{u\in S}\sum_{v\in\Gamma(v)\cap\bar{S}}2(\beta_{u,v}-\beta_{v,u})\\&-\delta^3\sum_{v\in S}\sum_{u\in\Gamma(v)\cap\bar{S}}\bar{\beta}_{v,u}-\delta^3\sum_{v\in S}\sum_{u\in\Gamma(v)\cap S}\bar{\beta}_{v,u}\\&-\delta^3\sum_{v\in\bar{S}}\sum_{u\in\Gamma(v)\cap S}\beta_{v,u}-\delta^3\sum_{v\in\bar{S}}\sum_{u\in\Gamma(v)\cap\bar{S}}\beta_{v,u}\Bigg)\\=n^{-1}\Bigg(&2\delta^2\sum_{\{u,v\}\in E}(\beta_{u,v}-\beta_{v,u})^2-\delta^3\sum_{\{u,v\}\in E}\eta_{\{u,v\}}\Bigg).\end{aligned}$$

Substituting $n=2m\delta$ concludes the proof. □

By symmetry, the expected value of $y_{b'}$ is the same as that of $y_{a'}$.

Lemma 22. *Define $y_{b'} = 0$ if the flip operation is not good. Then $\mathbf{E}[y_{b'}] = \mathbf{E}[y_{a'}]$.*

Proof. It follows from Lemma 12. □

5.3.3 Proof of Lemma 17

We have $\sigma_{G'}(S) - \sigma_G(S) = y_a + y_b + y_{a'} + y_{b'}$, where we set $y_{b'} = 0$ if the flip is not good. Then, from Lemmas 18, 19, 21 and 22,

$$\begin{aligned}
\mathbf{E}[\sigma_{G'}(S)] - \sigma_G(S) &= \mathbf{E}[y_a + y_b] + 2\,\mathbf{E}[y_{a'}] \\
&= 2\delta m^{-1} \sum_{\{u,v\}\in E_G(S,\bar{S})} (\beta_{u,v} - \beta_{v,u})^2 - \delta^2 m^{-1} \sum_{\{u,v\}\in E_G(S,\bar{S})} \zeta_{\{u,v\}} \\
&\quad + 2\delta m^{-1} \sum_{\{u,v\}\in E\setminus E_G(S,\bar{S})} (\beta_{u,v} - \beta_{v,u})^2 - \delta^2 m^{-1} \sum_{\{u,v\}\in E\setminus E_G(S,\bar{S})} \zeta_{\{u,v\}} \\
&\quad + 2\delta m^{-1} \sum_{\{u,v\}\in E} (\beta_{u,v} - \beta_{v,u})^2 - \delta^2 m^{-1} \sum_{\{u,v\}\in E} \eta_{\{u,v\}} \\
&= 4\delta m^{-1} \sum_{\{u,v\}\in E} (\beta_{u,v} - \beta_{v,u})^2 - \delta^2 m^{-1} \sum_{\{u,v\}\in E} \gamma_{G,\{u,v\}}(S).
\end{aligned}$$

5.4 Expected Cut Strain After a Flip in an (s,φ)-Expander

In this section, we apply Lemma 17 to show the following statement for (s,φ)-expanders.

Lemma 23. *Let $G = (V,E)$ be a d-regular (s,φ)-expander, where $\varphi \le 1/2$ and $d \ge 15/\varphi$, and let $\{S,\bar{S}\}$ be any cut such that $s \le |S| \le n/2$ and $\sigma_G(S) \le s/4$. If $G' = \mathrm{Flip}(G)$, then*

$$\mathbf{E}[\sigma_{G'}(S)] \ge \sigma_G(S) + \big((3\varphi/7)\cdot(\varphi/15 - \delta)^2 - 12\delta\big)\cdot \delta m^{-1} c_G(S).$$

Before we prove Lemma 23, we state a simple corollary of it and of Lemma 14.

Corollary 24. *For the same assumptions as in Lemma 23, the conditional expectation of $\sigma_{G'}(S)$ given the event $\mathcal{P}_S$ that the flip is pertinent to cut $\{S,\bar{S}\}$ is*

$$\mathbf{E}[\sigma_{G'}(S) \mid \mathcal{P}_S] \ge \sigma_G(S) + \big((\varphi/7)\cdot(\varphi/15 - \delta)^2 - 4\delta\big)\cdot\delta.$$

Note that if φ is a positive constant and d is at least a large enough constant, then the above inequality becomes $\mathbf{E}[\sigma_{G'}(S) \mid \mathcal{P}_S] \ge \sigma_G(S) + \varepsilon\cdot\delta$, for some positive constant ε.

Proof of Corollary 24. Since operations non-pertinent to $\{S,\bar{S}\}$ do not change the strain of the cut,

$$\mathbf{E}[\sigma_{G'}(S) - \sigma_G(S)] = \mathbf{E}[\sigma_{G'}(S) - \sigma_G(S) \mid \mathcal{P}_S]\cdot \Pr[\mathcal{P}_S].$$

Substituting the bounds on $\Pr[\mathcal{P}_S]$ and $\mathbf{E}[\sigma_{G'}(S) - \sigma_G(S)]$ from Lemmas 14 and 23, respectively, completes the proof. □

We divide the proof of Lemma 23 into three lemmas. The first bounds the sum of all $\gamma_{G,\{u,v\}}(S)$.

Lemma 25. *For any d-regular graph $G = (V,E)$ and cut $\{S,\bar{S}\}$,*

$$\sum_{\{u,v\}\in E} \gamma_{G,\{u,v\}}(S) \le 12 c_G(S).$$

Proof. Let $\{u,v\} \in E$. If $u \in S$ and $v \in \bar{S}$, then $\gamma_{G,\{u,v\}}(S) = \beta_{u,v} + \bar{\beta}_{v,u} + 4\beta_{u,v}\bar{\beta}_{v,u}/\rho_{u,v} \le 6$. If $u, v \in S$, then $\gamma_{G,\{u,v\}}(S) = \bar{\beta}_{u,v} + \bar{\beta}_{v,u} + 2(\beta_{u,v}\bar{\beta}_{v,u} + \bar{\beta}_{u,v}\beta_{v,u})/\rho_{u,v} \le 3\bar{\beta}_{u,v} + 3\bar{\beta}_{v,u} \le 3\bar{\alpha}_u + 3\bar{\alpha}_v$. Similarly, if $u, v \in \bar{S}$, then $\gamma_{G,\{u,v\}}(S) \le 3\alpha_u + 3\alpha_v$. Combining the three cases, we obtain

$$\sum_{\{u,v\}\in E} \gamma_{G,\{u,v\}}(S) \le 6c_G(S) + \sum_{u\in S}(3\bar{\alpha}_u \cdot d) + \sum_{v\in\bar{S}}(3\alpha_v \cdot d) = 12c_G(S). \qquad \square$$

The next lemma lower bounds the sum of all terms $(\beta_{u,v} - \beta_{v,u})^2$, when the cut size $c_G(S)$ is relatively large. In this case, it suffices to consider just the edges $\{u,v\}$ that cross the cut.

Lemma 26. *For any d-regular graph $G = (V,E)$ and cut $\{S, \bar{S}\}$, such that $c_G(S) \ge d\sigma_G(S)/(1-\epsilon)$ and $\delta \le \epsilon < 1$,*

$$\sum_{\{u,v\}\in E} (\beta_{u,v} - \beta_{v,u})^2 \ge (\epsilon - \delta)^2 \cdot c_G(S).$$

Proof. We have

$$\sum_{\{u,v\}\in E} (\beta_{u,v} - \beta_{v,u})^2 \ge \sum_{u\in S}\sum_{v\in\Gamma(u)\cap\bar{S}} (\beta_{u,v} - \beta_{v,u})^2 = \sum_{u\in S}\sum_{v\in\Gamma(u)\cap\bar{S}} (\alpha_u + \bar{\alpha}_v - (1-\delta))^2. \qquad (3)$$

Also,

$$\sum_{u\in S}\sum_{v\in\Gamma(u)\cap\bar{S}} (\alpha_u + \bar{\alpha}_v) = d\sigma_G(S) \le (1-\delta)\cdot c_G(S) = \sum_{u\in S}\sum_{v\in\Gamma(u)\cap\bar{S}} (1-\delta),$$

where the first equation holds because for each $u \in S$, α_u appears in $d \cdot \bar{\alpha}_u$ terms of the double sum, and for each $v \in \bar{S}$, $\bar{\alpha}_v$ appears in $d \cdot \alpha_v$ terms. It follows that the double sum in (3) is minimized when all the $c_G(S)$ many terms $(a_u + \bar{a}_v)$ are equal, thus

$$\sum_{\{u,v\}\in E} (\beta_{u,v} - \beta_{v,u})^2 \ge c_G(S) \cdot \left(\frac{d\sigma_G(S)}{c_G(S)} - (1-\delta)\right)^2 \ge c_G(S)\cdot(\epsilon-\delta)^2. \qquad \square$$

Next we lower bound the sum of all terms $(\beta_{u,v} - \beta_{v,u})^2$, when the cut size $c_G(S)$ is relatively small. This case is more involved as it does not suffice to consider just the crossing edges.

Lemma 27. *Let $G = (V,E)$ be a d-regular (s,φ)-expander, where $\varphi \le 1/2$ and $d \ge 15/\varphi$, and let $\{S, \bar{S}\}$ be a cut such that $s \le |S| \le n/2$ and $c_G(S) \le 7sd/27$. Then*

$$\sum_{\{u,v\}\in E} (\beta_{u,v} - \beta_{v,u})^2 \ge (3\varphi/28)\cdot(\varphi/15 - \delta)^2 \cdot c_G(S).$$

Proof. For real numbers $0 \le x \le y \le 1$, let $R_x = \{u \colon a_u \ge x\}$ and $D_{x,y} = R_x \setminus R_y$.

Claim 27.1. For any $0 \le x \le 1$,

$$|S| - \delta \cdot c_G(S)/(1-x) \le |R_x| \le |S| + \delta \cdot c_G(S)/x.$$

Proof of Claim 27.1. We have $|R_x| \ge |R_x \cap S| = |S| - |S \setminus R_x|$, and

$$\delta \cdot c_G(S) = \sum_{u\in S} \bar{a}_u \ge \sum_{u\in S\setminus R_x} \bar{a}_u \ge |S \setminus R_x| \cdot (1-x).$$

Thus, $|R_x| \ge |S| - \delta \cdot c_G(S)/(1-x)$. Similarly, $|R_x| \le |S| + |R_x \cap \bar{S}|$, and

$$\delta \cdot c_G(S) = \sum_{v\in\bar{S}} a_v \ge \sum_{v\in R_x\cap\bar{S}} a_v \ge |R_x \cap \bar{S}| \cdot x.$$

Thus $|R_x| \le |S| + \delta \cdot c_G(S)/x$. $\square$

Claim 27.2. For any $0 < a < b < 1$ and $0 < \epsilon < (b-a)/2$, there exists $x \in [a+\epsilon, b-\epsilon]$ such that

$$|E(D_{x-\epsilon,x}, D_{x,x+\epsilon})| \leq \frac{1}{\lfloor (b-a-\epsilon)/\epsilon \rfloor} \cdot \frac{1}{2}\left(\frac{1}{a} + \frac{1}{1-b}\right) c_G(S).$$

Proof of Claim 27.2. The volume of $D_{a,b}$ is at most $\left(\frac{1}{a} + \frac{1}{1-b}\right) c_G(S)$, and $|E(D_{a,b}, D_{a,b})|$ is at most half that volume. The claim follows by observing $|E(D_{a,b}, D_{a,b})| \geq \sum_{x=a+i\epsilon} |E(D_{x-\epsilon,x}, D_{x,x+\epsilon})|$, where the range of i in the summation is $1 \leq i \leq (b-a-\epsilon)/\epsilon$. □

Let $1/3 < x < 2/3$. Since $c_G(S) \leq 7sd/27$, Claim 27.1 gives $|S| - 7s/9 \leq |R_x| \leq |S| + 7s/9$. And since $s \leq |S| \leq n/2$, we get $2s/9 \leq |R_x| \leq n - 2s/9$. Then, by the fact G is an (s, φ)-expander,

$$c_G(R_x) \geq 2\varphi sd/9.$$

Fix the value of x such that it satisfies Claim 27.2, for $a = 1/3$, $b = 2/3$, and $\epsilon = \varphi/15$. Then,

$$|E(D_{x-\epsilon,x}, D_{x,x+\epsilon})| \leq \frac{3c_G(S)}{4/\varphi} \leq \frac{3 \cdot 7sd/27}{4/\varphi} = 7\varphi sd/36,$$

where in the first inequality we used $\varphi \leq 1/2$. It follows that at most $7\varphi sd/36$ of the $c_G(R_x) \geq 2\varphi sd/9$ edges $\{u, v\}$ that cross the cut $\{R_x, V \setminus R_x\}$ satisfy $|\alpha_u - \alpha_v| \leq \epsilon$. Thus $|\alpha_u - \alpha_v| > \epsilon$ for at least $2\varphi sd/9 - 7\varphi sd/36 = \varphi sd/36$ edges. And since $|\alpha_u - \alpha_v| > \epsilon$ implies $|\beta_{u,v} - \beta_{v,u}| > \epsilon - \delta$,

$$\sum_{\{u,v\} \in E} (\beta_{u,v} - \beta_{v,u})^2 \geq (\varphi sd/36) \cdot (\epsilon - \delta)^2 \geq (3\varphi \cdot c_G(S)/28) \cdot (\varphi/15 - \delta)^2. \qquad \square$$

Proof of Lemma 23. From Lemma 17,

$$\mathbf{E}[\sigma_{G'}(S)] = \sigma_G(S) + 4\delta m^{-1} \sum_{\{u,v\} \in E} (\beta_{u,v} - \beta_{v,u})^2 - \delta^2 m^{-1} \sum_{\{u,v\} \in E} \gamma_{G,\{u,v\}}(S).$$

We use Lemma 25 to upper bound the sum of $\gamma_{G,\{u,v\}}(S)$, and Lemmas 26 and 27 to lower bound the sum of $(\beta_{u,v} - \beta_{v,u})^2$ when $c_G(S) \geq (28d/27) \cdot \sigma_G(S)$ and $c_G(S) \leq (28d/27) \cdot \sigma_G(S) \leq 7sd/27$, respectively. □

6 Edge-Connectivity Analysis

In this section, we prove the follow lemma for an (n, d)-flip-chain, which states that after $r_c = O(nd \log^2 n)$ operations, the min-cut size $c(G_t)$ becomes at least $d/2$, and then remains at lease $d/2$ for a polynomial number of operations, w.h.p. The lemma requires that $d = \Omega(\log^2 n)$.

Lemma 28 (Edge-Connectivity Lower Bound)**.** *Let $\left(G_t = (V, E_t)\right)_{t \geq 0}$ be an (n, d)-flip-chain, and let $\tau_c = \min\{t \colon t \geq r_c,\, c(G_t) < d/2\}$, where $r_c = 292nd \ln^2 n$. For all $\ell > r_c$,*

$$\Pr[\tau_c \geq \ell] = 1 - O\left(n^{-3} d \log n + (3/5)^{d/(8 \log n)} \ell^2 n \log n\right).$$

6.1 Analysis of a Single Cut

We first study how the size of a single fixed cut $\{S, \bar{S}\}$ evolves. Recall that we replace subscript G_t in all notations by t, e.g., we write $c_t(S)$ and $\sigma_t(S)$ instead of $c_{G_t}(S)$ and $\sigma_{G_t}(S)$, respectively.

Lemma 29. *If $d \geq 171 \ln n$ then, for any cut $\{S, \bar{S}\}$, the time $\tau = \min\{t \colon c_t(S) \geq 3d/4\}$ satisfies*

$$\Pr[\tau \leq 292\, nd \ln n] \geq 1 - 3dn^{-4}.$$

Proof. We call the t-th operation, *flip*$_t$, *significant* if $\{a_t, b_t\} \in E_{t-1}(S, \bar{S})$.[8] Observe that non-significant operations do not change the *size* of cut $\{S, \bar{S}\}$. Let t_i denote the time of the ith significant operation for $i \geq 1$, and let $t_0 = 0$.

We will bound the number ρ of significant operations until time τ, and the time between successive significant operations. Let $\rho = \min\{i \colon c_{t_i}(S) \geq 3d/4\}$, thus $t_\rho = \tau$. For $i \geq 0$, let

$$X_i = \begin{cases} c_{t_i}(S) - i/2 & \text{if } 0 \leq i \leq \rho \\ X_\rho & \text{if } i > \rho. \end{cases}$$

From Lemma 16, it follows that $X_0, X_1, \ldots$ is a submartingale. Also, $-5/2 \leq X_{i+1} - X_i \leq 3/2$, for any $i \geq 0$. Then from Azuma's inequality (see, e.g., [28]), we obtain that for all $i \geq 1$ and $\lambda \geq 0$,

$$\Pr[X_i - X_0 > -\lambda] \geq 1 - e^{-2\lambda^2/(4^2 i)}.$$

Substituting the definition of X_i, we get

$$\Pr[\{c_{t_i} > i/2 - \lambda\} \cup \{i > \rho\}] \geq 1 - e^{-2\lambda^2/(4^2 i)}.$$

For $i \geq i^* \coloneqq 512 \ln n$ and $\lambda = \sqrt{4 \cdot 8 i \ln n}$, we have $i \geq 4\lambda$, and the above inequality gives

$$\Pr[\{c_{t_i}(S) > i/4\} \cup \{i > \rho\}] \geq 1 - n^{-4}.$$

For $i = 3d \geq 3 \cdot 171 \ln n > i^*$, the event $\{c_{t_i}(S) > i/4\} \cup \{i > \rho\}$ implies $\rho \leq 3d$. It follows

$$\begin{aligned} \Pr\left[\{\rho \leq 3d\} \cap \bigcap_{i^* \leq i < \rho} \{c_{t_i}(S) > i/4\}\right] &\geq \Pr\left[\bigcap_{i^* \leq i \leq 3d} \{c_{t_i}(S) > i/4\} \cup \{i > \rho\}\right] \\ &\geq 1 - (3d - i^* + 1) \cdot n^{-4}. \end{aligned}$$

The event $\bigcap_{i^* \leq i < \rho} \{c_{t_i}(S) > i/4\}$ above implies that if $i^* \leq i < \rho$, then $c_t(S) > i/4$ for $t_i \leq t < t_{i+1}$. Also, $c_t(S) \geq 1$ for $t < t_{i^*}$, since G_t is connected. We use these observations to bound the time between successive significant operations. For $i \geq 1$, let

$$Y_i = \begin{cases} t_i - t_{i-1} & \text{if } 1 \leq i \leq \rho \text{ and either } i \leq i^* \text{ or } c_{t_{i-1}}(S) > (i-1)/4 \\ 0 & \text{otherwise.} \end{cases}$$

Let $Y = \sum_{1 \leq i \leq 3d} Y_i$. It is easy to see that Y is dominated by the sum of $3d$ *independent* geometric random variables, $Y'_1, \ldots, Y'_{3d}$, where Y'_i has success probability $\frac{1}{m}$ if $i \leq i^*$ and $\frac{i-1}{4m}$ if $i > i^*$. Thus,

$$\mathbf{E}[Y] \leq i^* \cdot m + \sum_{i^* < i \leq 3d} \frac{4m}{i-1} \leq m\,(i^* + 4 \ln d) \leq 516 m \ln n =: \bar{y}.$$

[8] Significant operations are a subset of the operations pertinent to cut $\{S, \bar{S}\}$, described in Definition 13.

Then, by a tail bound on the sum of independent geometric random variables [17], for any $\lambda \geq 1$,

$$\Pr[Y < \lambda \bar{y}] \geq 1 - e^{-(1/m)\bar{y}(\lambda - 1 - \ln \lambda)}.$$

Setting $\lambda = 1.13$ yields

$$\Pr[Y < \lambda \bar{y}] \geq 1 - n^{-4}.$$

Finally, by a union bound,

$$\Pr\left[\{Y < \lambda \bar{y}\} \cap \{\rho \leq 3d\} \cap \bigcap_{i^* \leq i < \rho} \{c_{t_i}(S) > i/4\}\right] \geq 1 - 3dn^{-4}.$$

The event on the left implies $t_\rho \leq \lambda \bar{y} < 584m \ln n$. Since $t_\rho = \tau$, the proof is completed. □

Lemma 30. *For any cut $\{S, \bar{S}\}$ of initial size $c_0(S) \leq 3d/4$, and any even integer $0 < \lambda < c_0(S)$, the time $\tau = \min\{t \colon c_t(S) \leq c_0(S) - \lambda\}$ satisfies the next inequality, for all $\ell \geq 0$,*

$$\Pr[\tau > \ell] \geq 1 - \ell \cdot (3/5)^{\lambda/2}.$$

Proof. We call operation *flip*$_t$ *effective* if it changes the size of cut $\{S, \bar{S}\}$, i.e., $c_t(S) \neq c_{t-1}(S)$.[9] Each effective operation changes the cut size by ± 2. Let $k = c_0(S) \leq 3d/4$.

We only consider effective operations *flip*$_t$ for which it holds $k - \lambda < c_{t-1}(S) \leq k$. It follows from Lemma 16, that for each of those effective operations, the probability that the cut size increases by 2 is at least $p = \frac{2 - \delta k}{2}$, and the probability it decreases by 2 is at most $q = \frac{\delta k}{2}$, since $2p - 2q = 2(1 - \delta k)$. Then, the probability that the cut size decreases to $k - \lambda$, before it increases to $k + 2$, when the initial cut size is k, is bounded from below by the probability that a biased nearest-neighbor random walk on the integers hits the (left) point $l = 0$ before hitting the (right) point $r = \lambda/2 + 1$, when it starts from point $\lambda/2$, and in each step moves to the right or left with probability p and q, respectively. The above random walk probability is (see, e.g., [22])

$$\frac{1 - (q/p)^{r-l-1}}{1 - (q/p)^{r-l}} \geq 1 - (q/p)^{r-l-1} \geq 1 - (3/5)^{\lambda/2}.$$

Then, the probability than we need to repeat the random walk more than ℓ times before it hits l is at least $1 - \ell \cdot (3/5)^{\lambda/2}$. □

6.2 Analysis of Multiple Cuts

We combine the results for a single cut from the previous section, with the upper bound on the number of cuts in Lemma 8, to argue about multiple cuts simultaneously.

Lemma 31. *Let $1 \leq s < n/2$ and $0 < \lambda < k \leq 3d/4$, where λ is an even integer. If $d \geq 171 \ln n$ then the times $\tau = \min\{t \colon c(G_t, s) \leq k\}$ and $\tau' = \min\{t \colon t \geq r, c(G_t, 2s) \leq k - \lambda\}$, where $r = \lfloor 292nd \ln n \rfloor$, satisfy the following equation, for all $\ell \geq r$,*

$$\Pr\left[\{\tau' > \ell\} \cup \{\tau \leq \ell\}\right] = 1 - O(dn^{-3} + n\ell^2 \cdot (3/5)^{\lambda/2}).$$

[9] Effective operations are a subset of the set of significant operations, defined in the proof of Lemma 29.

Proof. For $t \geq 0$, let

$$A_t = \begin{cases} \mathcal{C}_0(2s, k) & \text{if } t = 0 \\ \mathcal{C}_t(2s, k) \setminus \mathcal{C}_t(2s, k-2) & \text{if } t \geq 1. \end{cases}$$

For any $S \subseteq V$ and $t \geq 0$, let

$$p_{S,t} = \min\{t' \colon t' \geq t,\, c_{t'}(S) \leq k - \lambda\}.$$

If $S \in A_0$, then from Lemmas 29 and 30 it follows

$$\Pr[p_{S,r} > \ell] = 1 - O\left(dn^{-4} + \ell \cdot (3/5)^{\lambda/2}\right).$$

Similarly, for any $t \geq 1$, if we are given $G_t = G$, then for every $S \in A_t$, Lemma 30 yields

$$\Pr[p_{S,t} > \ell \mid G_t = G] = 1 - O\left(\ell \cdot (3/5)^{\lambda/2}\right).$$

We can assume that $c(G_0, s) > k$ (otherwise $\tau = 0$, and $\Pr[\{\tau' > \ell\} \cup \{\tau < \ell\}] = 1$). Then, from Lemma 8, we have $|A_0| = |\mathcal{C}_0(2s, k)| \leq n$, and by a union bound

$$\Pr\left[\bigcap_{S \in A_0} \{p_{S,r} > \ell\}\right] = 1 - O\left(dn^{-3} + n\ell \cdot (3/5)^{\lambda/2}\right).$$

Similarly, for $t \geq 1$, if $c(G_t, s) > k$ then $|A_t| \leq |\mathcal{C}_t(2s, k)| \leq n$. And since $\tau > t$ implies $c(G_t, s) > k$

$$\Pr\left[\{\tau \leq t\} \cup \bigcap_{S \in A_t} \{p_{S,t} > \ell\}\right] = 1 - O\left(n\ell \cdot (3/5)^{\lambda/2}\right).$$

Then, by a union bound,

$$\Pr\left[\left(\bigcap_{S \in A_0} \{p_{S,r} > \ell\} \cap \bigcap_{S \in A_t, 1 \leq t \leq \ell} \{p_{S,t} > \ell\}\right) \cup \{\tau \leq \ell\}\right] = 1 - O\left(dn^{-3} + n\ell^2 \cdot (3/5)^{\lambda/2}\right).$$

We conclude the proof by observing that the event on the left side inside the parenthesis implies $\tau' > \ell$. Indeed, $\tau' > \ell$ if and only if $p_{S,r} > \ell$ for all sets S of size $1 \leq |S| \leq 2s$; and for any such set S, if $p_{S,r} \leq \ell$ then S must belong to at least one set A_t, for $0 \leq t \leq \ell$. □

Proof of Lemma 28. Let $i_{\max} = \lfloor \log(n/d) \rfloor$, $r = \lfloor 292nd \ln n \rfloor$, and let $\lambda \leq \frac{d}{4i_{\max}}$ be an even positive integer. For $0 \leq i \leq i_{\max}$, let $s_i = 2^i d$, $k_i = \lfloor 3d/4 \rfloor - \lambda i$, and define

$$\tau_i = \max\{t \colon t \geq ir,\, c(G_t, s_i) \leq k_i\}.$$

Note that $n/2 < s_{i_{\max}} \leq n$, and $k_{i_{\max}} \geq \lfloor 3d/4 \rfloor - d/4$. We apply Lemma 31, for $s = s_i$ and $k = k_i$, to obtain that for all $0 \leq i < i_{\max}$ and $\ell \geq (i+1) \cdot r$,

$$\Pr\left[\{\tau_{i+1} > \ell\} \cup \{\tau_i \leq \ell\}\right] = 1 - O\left(dn^{-3} + n\ell^2 \cdot (3/5)^{\lambda/2}\right).$$

It follows that, for any $\ell \geq i_{\max} r$,

$$\Pr\left[\tau_{i_{\max}} > \ell\right] = \Pr\left[\{\tau_{i_{\max}} > \ell\} \cup \{\tau_0 \leq \ell\}\right] \geq 1 - i_{\max} \cdot O\left(dn^{-3} + n\ell^2 \cdot (3/5)^{\lambda/2}\right),$$

where the first equation holds because $\Pr[\tau_0 \leq \ell] = 0$, by Lemma 9. To complete the proof, we use that $i_{\max} \leq \ln n$, we set λ to be the largest even integer $\lambda \leq \frac{d}{4 \ln n}$, and we observe that $\tau_{i_{\max}} > \ell$ implies $c_t(G_t) \geq k_{i_{\max}} + 1 \geq d/2$ for all $i_{\max} r \leq t \leq \ell$. □

7 Expansion Analysis

In this section we prove the following lemma for an (n,d)-flip-chain. Roughly speaking, it states that as long as the min-cut size is at least $d/2$, the graph conductance becomes $\phi(G_t) \geq 1/32$ after $r_e = O(nd\log n)$ flip operations, and then remains at lease $1/32$ for a polynomial number of operations, w.h.p. The lemma requires that $d = \Omega(\log n)$.

Lemma 32 (Expansion Lower Bound)**.** *Let $\left(G_t = (V, E_t)\right)_{t\geq 0}$ be an (n,d)-flip-chain, where $d \geq 26880\ln n/\nu_e^3$ and $\nu_e = (1/224)\cdot(1/480-\delta)^2 - 4\delta$. Let $\tau_h = \min\{t\colon c(G_t) < d/2\}$ and $\tau_e = \min\{t\colon t \geq r_e,\ \phi(G_t) < 1/32\}$, where $r_e = 48nd\log n/\nu_e$. For all $\ell \geq r_e$,*

$$\Pr\left[\{\tau_e > \ell\} \cup \{\tau_h \leq \ell\}\right] = 1 - O\left(\ell^2 \cdot e^{-\nu_e^2 d/2240}\right).$$

7.1 Analysis of a Single Cut

We first study how the strain of a single given cut evolves. The next three lemmas show that, roughly speaking, as long as the graph G_t remains an (s,φ)-expander, each the following statements holds with probability $1 - e^{-\Omega(sd)}$, for any given cut $\{S, \bar{S}\}$ with $s \leq |S| \leq n/2$: (1) the strain of the cut exceeds $s/4$ after at most $O(m)$ operation; (2) if the cut size is at least γsd, for a large enough constant γ, then the cut strain exceeds $s/4$ before the cut size drops to sd; and, (3) once the cut strain is at least $s/4$, it does not drop to $s/8$ in the next polynomial number of operations.

Lemma 33. *Let $1 \leq s \leq n/2$, $0 < \varphi \leq 1/2$, and $\nu = (\varphi/7)\cdot(\varphi/15-\delta)^2 - 4\delta$. If $\nu > 0$ then for any cut $\{S, \bar{S}\}$ with $s \leq |S| \leq n/2$, the time $\tau = \min\{t\colon \sigma_t(S) > s/4$ or $G_t \notin \mathbb{G}_{s,\varphi}\}$[10] satisfies*

$$\Pr\left[\tau \leq \frac{3m}{\varphi\nu}\right] \geq 1 - 2e^{-\nu sd/32}.$$

Proof. We will say that operation *flip*$_t$ is *pertinent* if it is pertinent to cut $\{S,\bar{S}\}$, i.e., if $\{a_t, b_t\} \in E_{t-1}(S,\bar{S})$, or $\{a_t, a'_t\} \in E_{t-1}(S,\bar{S})$, or if *flip*$_t$ is good and $\{b, b'\} \in E_{t-1}(S,\bar{S})$. Non-pertinent operations do not change the set of edges crossing the cut $\{S,\bar{S}\}$, thus neither the size nor the strain of the cut. Let t_i denote the time of the ith pertinent operation for $i \geq 1$, and let $t_0 = 0$.

We first bound the number of pertinent operations until time τ, similarly to the proof of Lemma 29. For $i \geq 0$, let

$$X_i = \begin{cases} \sigma_{t_i}(S) - i\cdot\nu\delta & \text{if } t_i \leq \tau \\ X_{i-1} & \text{if } t_i > \tau. \end{cases}$$

From Corollary 24, it follows that the sequence $X_0, X_1, \ldots$ is a submartingale.[11] Also, for any $i \geq 1$, $-4\delta - \nu\delta \leq X_i - X_{i-1} \leq 4\delta - \nu\delta$. Then, from Azuma's inequality [28], for all $i \geq 1$ and $\lambda \geq 0$,

$$\Pr\left[X_i - X_0 > -\lambda\delta\right] \geq 1 - e^{-\lambda^2/(32i)}.$$

Substituting the definition of X_i, we obtain

$$\Pr\left[\{\sigma_{t_i}(S) - \sigma_0(S) > i\nu\delta - \lambda\delta\} \cup \{t_i > \tau\}\right] \geq 1 - e^{-\lambda^2/(32i)}. \quad (4)$$

For $i = i^* \coloneqq \lceil 2sd/\nu \rceil$ and $\lambda = \sqrt{\nu sdi}$, we have $i\nu\delta - \lambda\delta \geq s/4$, and (4) gives

$$\Pr[t_{i^*} \geq \tau] \geq 1 - e^{-\nu sd/32}. \quad (5)$$

[10] Recall from Definition 7 that $\mathbb{G}_{s,\varphi}$ is the set of all (s,φ)-expander graphs.

[11] With respect to the sequence $Z_0, Z_1, \ldots$, where $Z_i = (G_{t_i}, \mathbb{1}_{\tau < t_{i+1}})$.

Next, we bound the time between pertinent operations. For $i \geq 1$, let

$$Y_i = \begin{cases} t_i - t_{i-1} & \text{if } t_{i-1} < \tau \\ 0 & \text{otherwise.} \end{cases}$$

Let $Y = \sum_{1 \leq i \leq i^*} Y_i$. The probability of the event $\mathcal{P}_t$ that *flip*$_t$ is pertinent is

$$\Pr[\mathcal{P}_t \mid G_{t-1}] \geq \Pr[\{a_t, b_t\} \in E_{t-1}(S, \bar{S}) \mid G_{t-1}] = c_{t-1}(S)/m.$$

Using this, and the fact that $c_t(S) \geq \varphi sd$ if $t < t_i$ and $t_{i-1} < \tau$, it is easy to see that Y is dominated by the sum of i^* i.i.d. geometric random variables with success probability $\varphi sd/m$. Applying a tail bound from [17], gives for any $\lambda \geq 1$,

$$\Pr\left[Y < \lambda \cdot \frac{i^* m}{\varphi sd}\right] \geq 1 - e^{-i^*(\lambda - 1 - \ln \lambda)}.$$

Setting $\lambda = 1.1$ and using that $i^* < 3sd/\nu$, we obtain $\Pr\left[Y < \frac{3m}{\varphi\nu}\right] \geq 1 - e^{-sd}$. Combining this and (5), by applying a union bound, yields $\Pr\left[\tau < \frac{3m}{\varphi\nu}\right] \geq 1 - 2e^{-\nu sd/32}$. □

Lemma 34. *Let $1 \leq s \leq n/2$, $0 < \varphi \leq 1/2$, and define ν as in Lemma 33. If $\nu > 0$ then for any cut $\{S, \bar{S}\}$ with $s \leq |S| \leq n/2$ and $c_0(S) \geq 5sd/\nu$, the times $\tau = \min\{t \colon \sigma_t(S) > s/4$ or $G_t \notin \mathbb{G}_{s,\varphi}\}$ and $\tau' = \min\{t \colon c_t(S) \leq sd\}$, satisfy*

$$\Pr\left[\tau < \tau'\right] \geq 1 - e^{-\nu sd/32}.$$

Proof. We saw in (5) that, for $i^* = \lceil 2sd/\nu \rceil$, the probability that the i^*th pertinent operation does not occur before time τ is

$$\Pr[t_{i^*} \geq \tau] \geq 1 - e^{-\nu sd/32}.$$

Since the size of cut $\{S, \bar{S}\}$ decreases by at most 2 in each pertinent operation, we have that for all $0 \leq i' \leq i^*$,

$$c_{t_{i'}}(S) \geq c_0(S) - 2i^* > sd.$$

Thus, $\tau' > t_{i^*}$ (w.pr. 1), and $\Pr[\tau' > t_{i^*} \geq \tau] \geq 1 - n^{-\nu sd/32}$. □

Lemma 35. *Let $1 \leq s \leq n/2$, $0 < \varphi \leq 1/2$, and define ν as in Lemma 33. If $\nu > 0$ and $d \geq 4225/(\nu s)$, then for any cut $\{S, \bar{S}\}$ with $s \leq |S| \leq n/2$ and $s/4 - 4\delta < \sigma_0(S) \leq s/4$, times $\tau_1 = \min\{t \colon \sigma_t(S) \leq s/8\}$ and $\tau_2 = \min\{t \colon G_t \notin \mathbb{G}_{s,\varphi}\}$ satisfy the next inequality, for all $\ell \geq 0$,*

$$\Pr\left[\{\tau_1 > \ell\} \cup \{\tau_2 \leq \ell\}\right] \geq 1 - \ell \cdot e^{-\nu^2 sd/1120}.$$

Proof. Let $\tau_3 = \min\{t \colon \sigma_t(S) > s/4\}$. We use the same submartingale sequence $X_0, X_1, \ldots$ as in the proof of Lemma 33, for

$$\tau = \min\{t \colon \sigma_t(S) > s/4 \text{ or } G_t \notin \mathbb{G}_{s,\varphi} \text{ or } \sigma_t(S) \leq s/8\} = \min\{\tau_1, \tau_2, \tau_3\}.$$

We will use (4), which we restate here for convenience,

$$\Pr\left[\{\sigma_{t_i}(S) - \sigma_0(S) > i\nu\delta - \lambda\delta\} \cup \{t_i > \tau\}\right] \geq 1 - e^{-\lambda^2/(32i)}.$$

For $i = i^* := \lfloor sd/32 \rfloor - 1$ and $\lambda = \nu \cdot \sqrt{sdi/35} \leq \nu sd/33$, we have

$$i\nu\delta - \lambda\delta \geq \nu s/32 - 2\nu\delta - \nu s/33 \geq 4\delta,$$

where the last inequality holds because $d \geq \frac{4225}{\nu s}$. Hence, for these values of i and λ, the event $\sigma_t(S) - \sigma_0(S) > \nu\delta i - \lambda\delta$ implies that $\sigma_{t_i}(S) > s/4$, and thus, that $t_i \geq \tau_3 \geq \tau$. Then, (4) yields

$$\Pr[t_{i^*} \geq \tau] \geq 1 - e^{-\nu^2 sd/1120}.$$

Moreover, since the cut strain decreases by at most 4δ in each pertinent operation, we have for all $0 \leq i' \leq i^*$,

$$\sigma_{t_{i'}}(S) \geq \sigma_0(S) - i' \cdot 4\delta > s/4 - 4\delta - i^* \cdot 4\delta \geq s/8.$$

Therefore, $t_{i^*} < \tau_1$ (w.pr. 1). It follows that

$$\Pr\left[\tau_1 > t_{i^*} \geq \tau = \min\{\tau_2, \tau_3\}\right] \geq 1 - e^{-\nu^2 sd/1120}.$$

I.e., with probability $1 - e^{-\nu^2 sd/1120}$, the cut strain exceeds $s/4$ or the graph becomes a non-(s, φ)-expander, before the cut strain decreases to $s/8$. The proof is completed by repeating the argument ℓ times, or until the graph is not an (s, φ)-expander. □

7.2 Analysis of Multiple Cuts

We combine the results for a single cut from the previous section, with the upper bound on the number of cuts from Theorem 10, to argue about multiple cuts simultaneously. We also use the simple relation between the strain and size of a cut from Lemma 4.

Lemma 36. *Let $1 \leq s \leq n/2$, $\varphi = 1/32$, $\varphi_0 = 1/2$, and define ν as in Lemma 33. If $\nu > 0$ and $d \geq 26880 \ln n/\nu^3$, then times $\tau = \min\{t\colon G_t \notin \mathbb{G}_{s,\varphi} \cup \mathbb{G}_{1,\varphi_0}\}$ and $\tau' = \min\{t\colon t \geq r, G_t \notin \mathbb{G}_{2s,\varphi}\}$, where $r = \lfloor \frac{3m}{\varphi\nu} \rfloor$, satisfy the next equation, for all $\ell \geq r$,*

$$\Pr\left[\{\tau' > \ell\} \cup \{\tau \leq \ell\}\right] = 1 - O\left(\ell^2 \cdot e^{-\nu^2 sd/2240}\right).$$

Proof. Let

$$A_0 = \{S\colon s \leq |S| \leq n/2, \sigma_0(S) \leq s/4, c_0(S) \leq k\},$$

where $k = \lfloor 6sd/\nu \rfloor$, and for $t \geq 1$, let

$$A_t = \{S\colon s \leq |S| \leq n/2, s/4 - 4\delta < \sigma_t(S) \leq s/4 \text{ or } k - 1 \leq c_t(S) \leq k\}.$$

For any $S \subseteq V$ and $t \geq 0$, let

$$p_{S,t} = \min\{t'\colon t' \geq t, \sigma_{t'}(S) \leq s/8, c_{t'}(S) \leq sd\}.$$

If $S \in A_0$, then from Lemmas 33 and 35 it follows that for all $\ell \geq r$,

$$\Pr[\{p_{S,r} > \ell\} \cup \{\tau \leq \ell\}] = 1 - O\left(\ell \cdot e^{-\nu^2 sd/1120}\right).$$

Similarly, for any $t \geq 1$, if we are given $G_t = G$, then for every $S \in A_t$,

$$\Pr[\{p_{S,t} > \ell\} \cup \{\tau \leq \ell\} \mid G_t = G] = 1 - O\left(\ell \cdot e^{-\nu^2 sd/1120}\right),$$

where the above equation follows from Lemma 35 if $s/4 - 4\delta < \sigma_t(S) \leq s/4$, and from Lemmas 34 and 35 if $k - 1 \leq c_t(S) \leq k$. Finally, for any $t \geq 0$, if $G_t \in \mathbb{G}_{1,\varphi_0}$ then Theorem 10 yields

$$|A_t| = O\left(n^{k/(\varphi_0 d)}\right) = O\left(n^{12s/\nu}\right).$$

Combining the above results, using a union bound, we obtain

$$\Pr\left[\left(\bigcap_{S\in A_0}\{p_{S,r}>\ell\}\cap\bigcap_{S\in A_t,1\le t\le\ell}\{p_{S,t}>\ell\}\right)\cup\{\tau\le\ell\}\right]=1-O\left(\ell^2\cdot e^{-\nu^2 sd/1120}\cdot n^{12s/\nu}\right)$$
$$=1-O\left(\ell^2\cdot e^{-\nu^2 sd/2240}\right),$$

where for the last equation we used that $d\ge 26880\ln n/\nu^3$. Observe that, for any set S of size $s\le|S|\le n/2$, if $p_{S,r}\le\ell$ then S must belong to at least one set A_t, for $0\le t\le\ell$. Therefore, the equation above implies

$$\Pr\left[\bigcap_{S\colon s\le|S|\le n/2}\{p_{S,r}>\ell\}\cup\{\tau\le\ell\}\right]=1-O\left(\ell^2\cdot e^{-\nu^2 sd/2240}\right).$$

Finally, from Lemma 4, if $p_{S,r}>\ell$ then for all $r\le t\le\ell$ we have $c_t(S)>sd/16$, and thus $c_t(S)>\varphi d\min\{2s,|S|\}$. Hence, the left side above is at most equal to $\Pr[\{\tau'>\ell\}\cup\{\tau\le\ell\}]$. □

Proof of Lemma 32. Let $\varphi=1/32$, $\varphi_0=1/2$, $r=\lfloor\frac{3m}{\varphi\nu}\rfloor$, and $i_{\max}=\lfloor\log n\rfloor$. Define the times $\tau_0=\min\{t\colon G_t\notin\mathbb{G}_{1,\varphi_0}\}$, and $\tau_i=\min\{t\colon t\ge ir, G_t\notin\mathbb{G}_{2^i,\varphi}\}$, for $1\le i\le i_{\max}$. We apply Lemma 36, for $s=2^i$, to obtain that for any $0\le i<i_{\max}$ and $\ell\ge(i+1)\cdot r$,

$$\Pr[\{\tau_{i+1}>\ell\}\cup\{\tau_i\le\ell\}\cup\{\tau_0\le\ell\}]=1-O\left(\ell^2\cdot e^{-\nu^2 sd/2240}\right).$$

It follows by a union bound that, for all $\ell\ge i_{\max}r$,

$$\Pr[\{\tau_{i_{\max}}>\ell\}\cup\{\tau_0\le\ell\}]=1-O\left(\ell^2\cdot e^{-\nu^2 d/2240}\right).$$

Observing that $\tau_{i_{\max}}\le\tau_e$, $\tau_0=\tau_h$, and $\nu=\nu_e$ for $\varphi=1/32$, completes the proof. □

8 Proof of Theorem 1

The theorem follows by combining Lemmas 28 and 32.

Acknowledgements

I would like to thank Péter L. Erdös for discussing with me some related previous works.
This work was partially supported by the French ANR Projects ByBloS (ANR-20-CE25-0002), DESCARTES (ANR-16-CE40-0023), and PAMELA (ANR-16-CE23-0016).

References

[1] Zeyuan Allen-Zhu, Aditya Bhaskara, Silvio Lattanzi, Vahab S. Mirrokni, and Lorenzo Orecchia. Expanders via local edge flips. In *Proc. 27th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 259–269, 2016. doi:10.1137/1.9781611974331.ch19.

[2] Georgios Amanatidis and Pieter Kleer. Rapid mixing of the switch Markov chain for strongly stable degree sequences. *Random Struct. Algorithms*, 57(3):637–657, 2020. doi:10.1002/rsa.20949.

[3] James Aspnes and Udi Wieder. The expansion and mixing time of skip graphs with applications. *Distributed Comput.*, 21(6):385–393, 2009. doi:10.1007/s00446-008-0071-3.

[4] Colin Cooper, Martin Dyer, and Catherine Greenhill. Corrigendum: Sampling regular graphs and a peer-to-peer network, 2012. arXiv:1203.6111.

[5] Colin Cooper, Martin E. Dyer, and Catherine S. Greenhill. Sampling regular graphs and a peer-to-peer network. *Comb. Probab. Comput.*, 16(4):557–593, 2007. doi:10.1017/S0963548306007978.

[6] Colin Cooper, Martin E. Dyer, Catherine S. Greenhill, and Andrew J. Handley. The flip Markov chain for connected regular graphs. *Discret. Appl. Math.*, 254:56–79, 2019. doi:10.1016/j.dam.2018.06.019.

[7] Colin Cooper, Martin E. Dyer, and Andrew J. Handley. The flip markov chain and a randomising P2P protocol. In *Proc. 28th Annual ACM Symposium on Principles of Distributed Computing, PODC*, pages 141–150, 2009. doi:10.1145/1582716.1582742.

[8] Martin E. Dyer, Catherine S. Greenhill, Pieter Kleer, James Ross, and Leen Stougie. Sampling hypergraphs with given degrees. *Discret. Math.*, 344(11):112566, 2021. doi:10.1016/j.disc.2021.112566.

[9] Péter L. Erdös, Catherine S. Greenhill, Tamás Róbert Mezei, István Miklós, Daniel Soltész, and Lajos Soukup. The mixing time of switch Markov chains: A unified approach. *Eur. J. Comb.*, 99:103421, 2022. doi:10.1016/j.ejc.2021.103421.

[10] Péter L. Erdős, Tamás Róbert Mezei, István Miklós, and Dániel Soltész. Efficiently sampling the realizations of bounded, irregular degree sequences of bipartite and directed graphs. *Plos one*, 13(8):e0201995, 2018. doi:10.1371/journal.pone.0201995.

[11] Tomás Feder, Adam Guetz, Milena Mihail, and Amin Saberi. A local switch Markov chain on given degree graphs with application in connectivity of peer-to-peer networks. In *Proc. 47th Annual IEEE Symposium on Foundations of Computer Science, FOCS*, pages 69–76, 2006. doi:10.1109/FOCS.2006.5.

[12] Christos Gkantsidis, Milena Mihail, and Amin Saberi. Random walks in peer-to-peer networks: Algorithms and evaluation. *Perform. Evaluation*, 63(3):241–263, 2006. doi:10.1016/j.peva.2005.01.002.

[13] Thorsten Götte, Kristian Hinnenthal, Christian Scheideler, and Julian Werthmann. Time-optimal construction of overlay networks. In *Proc. 40th ACM Symposium on Principles of Distributed Computing, PODC*, pages 457–468, 2021. doi:10.1145/3465084.3467932.

[14] Catherine S. Greenhill. A polynomial bound on the mixing time of a Markov chain for sampling regular directed graphs. *Electron. J. Comb.*, 18(1):P234, 2011. doi:10.37236/721.

[15] Catherine S. Greenhill and Matteo Sfragara. The switch Markov chain for sampling irregular graphs and digraphs. *Theor. Comput. Sci.*, 719:1–20, 2018. doi:10.1016/j.tcs.2017.11.010.

[16] Riko Jacob, Andréa W. Richa, Christian Scheideler, Stefan Schmid, and Hanjo Täubig. SKIP$^+$: A self-stabilizing skip graph. *J. ACM*, 61(6):36:1–36:26, 2014. doi:10.1145/2629695.

[17] Svante Janson. Tail bounds for sums of geometric and exponential variables. *Statistics & Probability Letters*, 135:1–6, 2018. doi:10.1016/j.spl.2017.11.017.

[18] Márk Jelasity, Spyros Voulgaris, Rachid Guerraoui, Anne-Marie Kermarrec, and Maarten van Steen. Gossip-based peer sampling. *ACM Trans. Comput. Syst.*, 25(3):8, 2007. doi:10.1145/1275517.1275520.

[19] Ravi Kannan, Prasad Tetali, and Santosh S. Vempala. Simple Markov-chain algorithms for generating bipartite graphs and tournaments. *Random Struct. Algorithms*, 14(4):293–308, 1999. doi:10.1002/(SICI)1098-2418(199907)14:4<293::AID-RSA1>3.0.CO;2-G.

[20] David R. Karger. Minimum cuts in near-linear time. *J. ACM*, 47(1):46–76, 2000. doi:10.1145/331605.331608.

[21] Ching Law and Kai-Yeung Siu. Distributed construction of random expander networks. In *Proc. 22nd Annual Joint Conference of the IEEE Computer and Communications Societies, INFOCOM*, pages 2133–2143, 2003. doi:10.1109/INFCOM.2003.1209234.

[22] David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov chains and mixing times.* American Mathematical Society, 2006.

[23] Peter Mahlmann. *Peer-to-peer networks based on random graphs.* PhD thesis, University of Paderborn, 2010. URL: https://d-nb.info/1003757057/34.

[24] Peter Mahlmann and Christian Schindelhauer. Peer-to-peer networks based on random transformations of connected regular undirected graphs. In *Proc. 17th Annual ACM Symposium on Parallelism in Algorithms and Architectures, SPAA*, pages 155–164, 2005. doi:10.1145/1073970.1073992.

[25] Laurent Massoulié and Rémi Varloot. Rapid mixing of dynamic graphs with local evolution rules. *IEEE Trans. Netw. Sci. Eng.*, 7(3):1344–1354, 2020. doi:10.1109/TNSE.2019.2926887.

[26] Brendan D. McKay. Subgraphs of random graphs with specified degrees. *Congressus Numerantium*, 33:213–223, 1981.

[27] István Miklós, Péter L. Erdös, and Lajos Soukup. Towards random uniform sampling of bipartite graphs with given degree sequence. *Electron. J. Comb.*, 20(1):P16, 2013. doi:10.37236/3028.

[28] Michael Mitzenmacher and Eli Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis.* Cambridge University Press, 2005. doi:10.1017/CBO9780511813603.

[29] Scott Oaks and Li Gong. *JXTA in a Nutshell.* O'Reilly & Associates, Inc., USA, 2002.

[30] Gopal Pandurangan, Prabhakar Raghavan, and Eli Upfal. Building low-diameter P2P networks. In *Proc. 42nd Annual Symposium on Foundations of Computer Science, FOCS*, pages 492–499, 2001. doi:10.1109/SFCS.2001.959925.

[31] Konstantin Tikhomirov and Pierre Youssef. Sharp Poincaré and log-Sobolev inequalities for the switch chain on regular bipartite graphs, 2020. arXiv:2007.02729.

[32] Nicholas C. Wormald. *Models of random regular graphs*, pages 239—298. London Mathematical Society Lecture Note Series. Cambridge University Press, 1999. doi:10.1017/CBO9780511721335.010.