Electronic Colloquium on Computational Complexity, Report No. 101 (2011)

Testing Conductance in General Graphs [*]

Angsheng Li[†] Yicheng Pan [†] Pan Peng[†]

July 27, 2011

Abstract

In this paper, we study the problem of testing the conductance of a given graph in the general graph model. Given distance parameter ε and any constant $\sigma > 0$, there exists a tester whose running time is $\mathcal{O}(\frac{m^{(1+\sigma)/2}\cdot\log n\cdot\log\frac{1}{\varepsilon}}{\varepsilon\cdot\Phi^2})$, where n is the number of vertices and m is the number of edges of the input graph. With probability at least $2/3$, the tester accepts all graphs of conductance at least Φ, and rejects any graph that is ε-far from any graph of conductance at least α' for $\alpha' = \Omega(\Phi^2)$. This result matches the best testing algorithm for the bounded degree graph model in [5].

Our main technical contribution is the *non-uniform Zig-Zag product*, which generalizes the standard Zig-Zag product given by Reingold et. al. [9] to the unregular case. It converts any graph to a regular one and keeps (roughly) the size and conductance, by choosing a proper Zig-Zag graph sequence. This makes it easy to test the conductance of the given graph on the new one. The analysis and applications of non-uniform Zig-Zag product may be independently interesting.

1 Introduction

Given a graph $G = (V, E)$ on vertex set V and edge set E, let $S \subseteq V$ be a vertex subset. The *volume* of S is defined to be the summation of the degrees of the vertices in S, denoted by $\mathrm{vol}_G(S) = \sum_{v\in S}\deg(v)$. Given a cut $(S, \overline{S})$, where $\overline{S}$ denotes the complement of S in V, the *conductance of the cut* is defined to be

$$\mathrm{cond}_G(S) = \frac{E(S,\overline{S})}{\min\{\mathrm{vol}(S),\mathrm{vol}(\overline{S})\}},$$

[*]The research is partially supported by NSFC distinguished young investigator award number 60325206, and its matching fund from the Hundred-Talent Program of the Chinese Academy of Sciences. All authors are partially supported by the Grand Project "Network Algorithms and Digital Information" of the Institute of software, Chinese Academy of Sciences.

[†]State Key Laboratory of Computer Science, Institute of Software, Chinese Academy of Sciences. {angsheng,yicheng,pengpan}@ios.ac.cn. Corresponding author: Yicheng Pan

ISSN 1433-8092

where $E(S, \overline{S})$ is the number of edges crossing the cut. We also write $\mathrm{vol}(S)$ and $\mathrm{cond}(S)$ for abbreviation if G is clear from context. The *conductance of graph G* is the minimum conductance among all the cuts of G. That is,

$$\mathrm{cond}(G) = \min_S \frac{E(S, \overline{S})}{\min\{\mathrm{vol}(S), \mathrm{vol}(\overline{S})\}}.$$

In general graphs, such as networks, the conductance is often used as a criterion of small communities. A cut $(S, \overline{S})$ which has a small conductance implies a vertex subset S such that it is (relatively) dense inside S and (relatively) sparse on its boundary. In the special case of bounded degree graphs, whose maximum degree of vertices is upper bounded by a constant d, such criterion is always given by the edge expansion. In this case, both the vertex and edge expansions are bounded by a constant times (depending on d) the conductance. Testing expansion (essentially testing conductance) in bounded degree model has been studied for a long time, but people have no idea for the case of general graphs. In this paper, we investigate the this problem.

Firstly, we briefly introduce the study for the bounded degree case.

1.1 Testing conductance in bounded degree model

Graph property testing is the task to test whether a graph has a given property or far away from having it by a sublinear time randomized algorithm, called *tester*. The tester is given an oracle access to a suitable representation of the graph. It requires giving right answers with probability at least $\frac{2}{3}$. Different graph models have different representation. For example, in dense graph model, an n-vertex graph is usually represented by its adjacency matrix of size $n \times n$. For bounded degree d graph model, an n-vertex graph is stored in an $[n] \times d$ matrix, where $[n]$ denotes the integer set $\{1, \ldots, n\}$. The j-th component in row i represents the label of the j-th neighbor of vertex i. So it requires $d \cdot n \log n$ bits. In these two models, we say that a graph is ε-far from another one if there are at least ε fraction of components in their matrices are different despite of labeling.

The problem of testing expansion in bounded degree model was first formulated by Goldreich and Ron [3] in 2000. They gave a tester with analysis depending on an unproven combinatorial conjecture. In 2007, using combinatorial techniques, Czumaj and Sohler [2] proposed a tester for vertex expansion. They showed that, given parameters $\alpha, \varepsilon > 0$, the tester accepts all graphs with vertex expansion larger than α, and rejects all graphs that are ε-far from having vertex expansion less than $\alpha' = \Theta(\frac{\alpha^2}{d^2 \log n})$. Recently, using algebraic argument based on the idea of Goldreich and Ron in [3], Kale and Seshadhri [5] as well as Nachmias and Shapira [7] improved α' to $\Theta(\alpha^2)$ for both vertex and edge expansions. The constant in Θ depends on d and the query complexity is $\frac{1}{\alpha^2} \cdot \tilde{\mathcal{O}}(n^{\frac{1}{2}+\mu} \cdot \frac{1}{\varepsilon})$ [1] for any small constant $\mu > 0$. Because of the Cheeger obstacle, the α' is hard to improve. However, the query complexity is almost touching the lower bound $\Omega(\sqrt{n})$ [4] [2].

The subject they discuss for testing expansion is in fact to test conductance since expansion is bounded by conductance in this model. The main idea is based on random walks by the following intuition. If the graph has large conductance, then random walks starting from any vertex mix very fast, and they collide with each other with low probability. Otherwise, once starting from a node in a small set with small conductance, the random walks cannot go out of it easily, and then the random walks starting from this node will collide with high probability.

1.2 From bounded degree to general graphs – the non-uniform Zig-Zag product

It seems that the idea for bounded degree model can be used directly in general graphs. However, there is an essential difference: in the bounded degree model, there is no vertex has high degree, so that all vertices have almost the same status. For the lazy random walk defined in [5], we know that the stationary distribution is uniform, and the collision probability is a good estimate of conductance. In general graphs, such random walk requires too much time since the largest degree of all vertices may be as large as $\Theta(n)$. For other kinds of random walks, we cannot compute the probability distribution easily since we have no idea of the degree distribution (we cannot query most of them). So it is hard to estimate the conductance simply by collision probability. For example, suppose that we are given a star, in which the central node has degree $n-1$ and the surrounding $n-1$ nodes have degree one for each. The random walks on it collide on the central vertex with very high probability, but in fact, the star has large conductance.

How to decentralize the impact of such heavy vertices? We introduce a new "coding" technique for graphs, called *non-uniform Zig-Zag product*. The standard Zig-Zag product was formulated by Reingold, Vadhan and Wigderson [9] in 2000. It yields simple constructions of constant-degree expanders of arbitrary size and plays a central role in the proof of Undirected Connectivity in **L** [8]. It takes a product of a large *regular* graph with a *single* small graph by blowing each vertex in the large graph up to a "cloud" of a small number of vertices, and the resulting graph inherits roughly its size from the large one, its degree from the small one, and its expansion property from both. Here, the non-uniform Zig-Zag produnct generalizes the large graph in the standard Zig-Zag product to *unregular* one. It

[1] $\tilde{\mathcal{O}}(n)$ denotes $\mathcal{O}(n \log n)$.

[2] $\Omega(\frac{1}{\varepsilon})$ is also a lower bound, since to query on the corrupted segment of the input with high probability, $\Omega(\frac{1}{\varepsilon})$ times of query is necessary.

also generalizes the small graph to a *sequence* of small regular graphs, each of which having size d corresponds to vertices in the large graph having degree d. That is the reason why we call our product "non-uniform". We show that the resulting graph also inherits its size from the large graph, its degree from the small graph sequence, and its conductance property from both. Our main interests focus on the analysis of the second eigenvalue, which is closely related to conductance by Cheeger Inequality.

Suppose that $\mathcal{H}$ is a d-regular graph sequence. The Zig-Zag product of G and $\mathcal{H}$, denoted by $\widehat{G} = G \textcircled{z} \mathcal{H}$, satisfies that

(1) The size of $\widehat{G}$ is a linear expansion of the size of G (in proper representation).

(2) $\widehat{G}$ is d^2-regular.

(3) The second eigenvalue of $\widehat{G}$ is lower bounded by some function of eigenvalues of G and graphs in $\mathcal{H}$. (The Separation Lemma in the next subsection.)

The non-uniform Zig-Zag product converts any graph into a regular one, and keeps the conductance largely by choosing proper $\mathcal{H}$. This allows us to design a conductance tester for general graphs.

1.3 Outlines of our analysis

To test the conductance of G, we actually test the conductance of $\widehat{G}$. An important notion is that, under the restriction of query complexity on G, $\widehat{G}$ is just imaginary rather than constructed. Since $\widehat{G}$ is d^2-regular, the tester for bounded degree model works on it. Recall that the tester can distinguish the case of "having conductance at least Φ" from the case of "ε-far from having conductance at least $c\Phi^2$" for some constant c (See [5], Theorem 1.1). We show that based on this tester, we can distinguish for G the case of "having conductance at least Φ" from the case of "ε-far from having conductance at least $\Omega(\Phi^2)$". Concretely, we prove the Separation Lemma as follows.

Lemma 1. *(Separation Lemma) Let Φ and Φ' be parameters in $[0, 1]$.*

(1) If $\operatorname{cond}(G) \geq \Phi$, then $\operatorname{cond}(\widehat{G}) \geq \lambda_{\widehat{G}}/2 \geq c' \cdot \Phi^2$ for some constant c' (by choosing a proper $\mathcal{H}$), where $\lambda_{\widehat{G}}$ is the second eigenvalue of the Laplacian of $\widehat{G}$.

(2) If $\operatorname{cond}(G) \leq \Phi'$, then an easy observation implies that $\operatorname{cond}(\widehat{G}) \leq \Phi'$.

The proof of (2) is given in Section 3 and (1) in Section 4. In fact, for (1), what we really want is $\lambda_{\widehat{G}}/2 \geq c' \cdot \Phi^2$. In the analysis of the tester for bounded degree model in [5], the gap between Φ and $\Omega(\Phi^2)$ stems from the process that the condition $\mathrm{cond}(G) \geq \Phi$ is first converted to $\lambda_G \geq \Phi^2/2$ by the Cheeger inequality. Here we show that $\lambda_{\widehat{G}} = \Omega(\Phi^2)$ directly and thus do not need to apply the Cheeger inequality once more. This allows us to keep the gap Φ versus $\Omega(\Phi^2)$ as [5] concludes, rather than a gap like Φ versus $\Omega(\Phi^4)$.

However, it is not enough yet. What we really want to show is that "if G is ε-far from having conductance at least $\Omega(\Phi^2)$, then the rejection probability is at least $2/3$". We show the contrapositive: if the rejection probability is too small, then there is a graph G' that is ε-close to G and has conductance at least $\Omega(\Phi^2)$. We show the existence of such G' by a patch-up algorithm, which changes G to G' by not too many modifications. Concretely, we show that if the rejection probability is small, then there exist few "weak" nodes in $\widehat{G}$, from which the random walks collide with each other with high probability. It implies that there are few "bad" nodes in G. Otherwise, they determine a relatively balanced cut with small conductance. Then the patch-up algorithm dealing with such bad nodes does not need to modify G a lot. We conclude our main result as follows.

Theorem 1. *(Main Theorem) Given any conductance parameter $0 \leq \Phi \leq 1$, distance parameter ε and any constant $\sigma > 0$, there exists a tester $\mathcal{T}$ whose running time is $\mathcal{O}(\frac{m^{(1+\sigma)/2} \cdot \log n \cdot \log \frac{1}{\varepsilon}}{\varepsilon \cdot \Phi^2})$, where n is the number of vertices and m is the number of edges of the input graph. With probability at least $2/3$, $\mathcal{T}$ accepts all graphs of conductance at least Φ, and rejects any graph that is ε-far from any graph of conductance at least $\Omega(\Phi^2)$.*

This result matches the best tester for conductance in the bounded degree model given by Kale and Seshadhri [5]. It seems that testing conductance in general graphs is no harder than testing it in bounded degree ones.

2 Preliminaries

In this section, we introduce some backgrounds of graph representation and spectral graph facts that we need in our paper.

2.1 Graph representation in the general graph model

We have mentioned in the above section that the dense and bounded degree models have canonical representation, using adjacency and incidence matrices, respectively. In this paper, we use the incidence representation for general graphs with slight modifications. For a graph $G = (V, E)$, let $|V| = n$ and $|E| = m$. For each vertex v, let d_v be the degree of v. The incidence representation is a list without a uniform yardstick as the bounded degree model does. Usually, the (u, i)-th component represents the index (among nodes in G) of the i-th neighbor of vertex u. Since each index needs $\log n$ bits, we need $\mathcal{O}(m \log n)$ bits to represent G.

For our application, we give a slight modification. Firstly, at the beginning of each row, we store the degree of the corresponding vertex. This allows us to know how many random bits we need when we want to choose a random neighbor. Secondly, in the (u, i)-th component for each $u \in V$ and $i \in [d_u]$, besides the index of the i-th neighbor of u, denoted by v, we store the index (among the neighbors of v) of u. That is, if u is the j-th neighbor of v, then we put j in the (u, i)-th component with the index of v together. It is easy to verify that the number of bits needed by this format is also $\mathcal{O}(m \log n)$, and such modification is reasonable.

So we have two kinds of query. The first is *degree query*. For each vertex u, we can query the list for d_u. The second is *neighbor query*. For each vertex u and index $i \in [d_u]$, we can query the (u, i)-th component for the i-th neighbor of u, denoted by v, and the index of u among v's neighbors. We define the relative distance between the graphs $G = (V, E)$ and $G' = (V, E')$ to be $\frac{|E \triangle E'|}{\max\{|E|, |E'|\}}$, where $\triangle$ refers to the symmetric difference.

Such representation implies a useful map in terms of which, it is convenient for us to define the non-uniform Zig-Zag product.

Definition 1. *(Rotation Map) For an undirected graph $G = (V, E)$, for $u, v \in V$, $i \in [d_u]$ and $j \in [d_v]$, the rotation map $Rot_G(u, i) = (v, j)$ if the i-th neighbor of u is v and the j-th neighbor of v is u.*

The difference of our rotation map from that defined in [9], Def. 2.1, is that our graph G does not need to be regular. It keeps track of the edge traversed to get from u to v, and $\text{Rot}_G(u, i) = (v, j)$ if and only if $\text{Rot}_G(v, j) = (u, i)$. So Rot_G is an exchange permutation satisfying that $\text{Rot}_G \circ \text{Rot}_G$ is identity.

Our graph representation *locally* implies the rotation map. That is, we only have to query the (u, i)-th component to get $(v, j) = \text{Rot}_G(u, i)$, instead of querying the v-th row to find j. In the construction of the non-uniform Zig-Zag product, we always suppose that we know the rotation map of graphs.

2.2 Some facts of spectral graphs

Let A be the adjacency matrix of graph G and D be the diagonal matrix with the (u, u)-th entry having value d_u. The *Laplacian* of G is defined to be

$$\mathcal{L} = I - D^{-\frac{1}{2}} A D^{-\frac{1}{2}},$$

where I is the identity matrix. $\mathcal{L}$ is symmetric. Its *spectrum* (or the *spectrum of* G) is defined to be the set of n eigenvalues of $\mathcal{L}$, denoted by $\{\lambda_1, \lambda_2, \cdots, \lambda_n\}$. It is well known that all the eigenvalues are in the interval $[0, 2]$ and the first eigenvalue $\lambda_1 = 0$. Suppose that $0 = \lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n \leq 2$, where λ_i is called the i-th eigenvalue of $\mathcal{L}$. We denote λ_2 by λ_G. Let f be a function mapping from V to $\mathbb{R}$. We have

Proposition 1.

$$\lambda_G = \inf_{f \perp D\vec{1}} \frac{\sum\limits_{u \sim v} (f(u) - f(v))^2}{\sum\limits_{v} f(v)^2 d_v},$$

where $\vec{1}$ denotes the all 1 vector and $\sum\limits_{u \sim v}$ denotes the sum over all unordered adjacent pairs $\{u, v\}$.

Let $M = I - \mathcal{L} = D^{-\frac{1}{2}} A D^{-\frac{1}{2}}$. The eigenvalues of M are $1 - \lambda_i$, $i = 1, \ldots, n$. To compute λ_G, sometimes it is convenient to compute the second largest eigenvalue of M, which is

$$1 - \lambda_G = \sup_{f \perp D\vec{1}} \frac{\sum\limits_{u \sim v} 2f(u)f(v)}{\sum\limits_{v} f(v)^2 d_v}. \tag{1}$$

For the relationship of $\mathrm{cond}(G)$ and λ_G, we have the Cheeger inequality.

Proposition 2. *(Cheeger Inequality)*

$$\frac{\mathrm{cond}(G)^2}{2} < \lambda_G \leq 2 \cdot \mathrm{cond}(G)$$

The conductance of G is bounded between $\lambda_G/2$ and $\sqrt{2\lambda_G}$ and thus λ_G is a measure of $\mathrm{cond}(G)$.

When G is d-regular, the matrix M is the normalized adjacency matrix $D^{-1}A$. The normalized eigenvector of λ_1 is $\vec{1}/\sqrt{n}$. Then we have

$$\lambda_G = \min_{\alpha \perp \vec{1}} \frac{< \alpha, \mathcal{L}\alpha >}{< \alpha, \alpha >} = 1 - \max_{\alpha \perp \vec{1}} \frac{< \alpha, M\alpha >}{< \alpha, \alpha >}, \tag{2}$$

and the second largest eigenvalue (in absolute value) of M, denoted by η_G, is given by

$$\eta_G = \max_{\alpha \perp \vec{1}} \frac{| < \alpha, M\alpha > |}{< \alpha, \alpha >} = \max_{\alpha \perp \vec{1}} \frac{\| M\alpha \|}{\| \alpha \|}. \tag{3}$$

$< \cdot, \cdot >$ denotes the standard inner product of vectors in $\mathbb{R}^n$ and $\| \cdot \|$ is the 2-norm. It is easily observed that $\eta_G = \max\{|1 - \lambda_G|, |1 - \lambda_n|\}$.

3 Non-uniform Zig-Zag product

In this section, we define the non-uniform Zig-Zag product formally and prove some basic properties. The non-uniform Zig-Zag product transforms an arbitrary graph G to a regular graph $\widehat{G}$ with size linear to $|E|$, and $\mathrm{cond}(\widehat{G})$ is properly bounded by $\mathrm{cond}(G)$.

Let $\mathcal{H} = \{H_1, H_2, \ldots\}$ be a d-regular graph family (multi-edges and self-loops are permitted). For each i, H_i has i vertices. We call such $\mathcal{H}$ a *Zig-Zag sequence*. For a vertex $u \in V$, let H_u denote the graph H_{d_u}, which is the graph in $\mathcal{H}$ that has size d_u. The non-uniform Zig-Zag product $G \textcircled{z} \mathcal{H}$ is a d^2-regular graph, in which every vertex u from G is blown up to be a "cloud" of d_u nodes, denoted by $(u, 1), \ldots, (u, d_u)$. The rotation map of $G \textcircled{z} \mathcal{H}$ is formally defined as follows.

Definition 2. *(Definition of Rot$_{G \textcircled{z} \mathcal{H}}$) For $u \in V$, $a \in [d_u]$ and $i, j \in [d]$,*

(1) Let (a', i') =Rot$_{H_u}(a, i)$.

(2) Let (v, b') =Rot$_G(u, a')$.

(3) Let (b, j') =Rot$_{H_v}(b', j)$.

Define Rot$_{G \textcircled{z} \mathcal{H}}((u, a), (i, j)) = ((v, b), (j', i'))$.

The intuition of this product is similar with that for the standard one defined by Reingold et. al. [9]. The three steps above correspond to three steps of walks in H_u, G and H_v, respectively. For some $(u, a) \in V \times [d_u]$, if we randomly choose $(i, j) \in [d] \times [d]$, which is one step of a random walk in $G \textcircled{z} \mathcal{H}$, then it is observed that the "Zig" step (1) represents a random step in H_u, to decide which cloud it is going into. The "long" step (2) represents a step crossing the edge (u, v) in G and the "Zag" step (3) is a random step in H_v. The important intuition for our application is that, when we choose each H_u in $\mathcal{H}$ to be an expander such that the "long" step decided by a' is as uniform as possible, it makes the random walk in $G \textcircled{z} \mathcal{H}$ look much like a random walk in G. The conductance of G will be hidden in $G \textcircled{z} \mathcal{H}$ very well. The proof of Theorem 3 follows this intuition closely.

Let $\widehat{G} = G \textcircled{z} \mathcal{H}$. For the relationship of $\mathrm{cond}(G)$ and $\mathrm{cond}(\widehat{G})$, the easy direction stated in Lemma 1 (2) is given as follows.

Lemma 2. *(Upper Bound Lemma)*

$$\mathrm{cond}(\widehat{G}) \leq \mathrm{cond}(G).$$

Proof. By definition, it is easy to verify the following two facts:

Fact 1. *For a vertex $u \in V$ in G, the volume of the corresponding cloud u in $\widehat{G}$ is exactly $d^2 \cdot d_u$.*

Fact 2. *For an edge $e = \{u, v\}$ in G, the number of edges contributed by e between the clouds u and v in $\widehat{G}$ is exactly d^2.*

So for any set $S \subseteq V$, we define $\widehat{S}$ to be the set of all nodes in $\widehat{G}$, each of which is in the cloud associated with some vertex in S. Then we know that $\mathrm{vol}(\widehat{S}) = d^2 \cdot \mathrm{vol}(S)$, $\mathrm{vol}(\neg\widehat{S}) = d^2 \cdot \mathrm{vol}(\overline{S})$ and $E(\widehat{S}, \neg\widehat{S}) = d^2 \cdot E(S, \overline{S})$, where $\neg\widehat{S}$ denotes the complement of $\widehat{S}$. The conductance of the cut $(\widehat{S}, \neg\widehat{S})$ equals the conductance of $(S, \overline{S})$. Since $\mathrm{cond}(\widehat{G})$ is the minimum conductance among all cuts in $\widehat{G}$, the proposition follows. $\square$

Then we show the hard direction, i.e., (1) of Lemma 1.

4 Lower bound for cond$(G \text{ⓩ} \mathcal{H})$

In this section, we show the following key lemma.

Lemma 3. *(Lower Bound Lemma) Let $G = (V, E)$ be a graph with $\mathrm{cond}(G) \geq \Phi$ for some Φ in $[0, 1]$, and $\mathcal{H}$ be a d-regular Zig-Zag sequence. Let $\eta = \max_{H \in \mathcal{H}}\{\eta_H\}$ and $\widehat{G} = G \text{ⓩ} \mathcal{H}$. Then $\mathrm{cond}(\widehat{G}) \geq \frac{\lambda_{\widehat{G}}}{2} \geq \frac{1}{8}(1 - \eta^2)\Phi^2$.*

Proof. Let A_G, $A_{\widehat{G}}$ be the adjacency matrices of G and $\widehat{G}$ respectively. Note that $A_{\widehat{G}}$ has size $\mathrm{vol}(G) \times \mathrm{vol}(G)$ and $\widehat{G}$ is d^2-regular. Its normalized adjacency matrix, denoted by M, is $\frac{1}{d^2} A_{\widehat{G}}$. By the Cheeger inequality, since $\mathrm{cond}(G) \geq \Phi$, we know that $\lambda_G > \frac{\mathrm{cond}(G)^2}{2} \geq \frac{\Phi^2}{2}$. Since $\mathrm{cond}(\widehat{G}) \geq \frac{1}{2}\lambda_{\widehat{G}}$, our following task is to relate $\lambda_{\widehat{G}}$ to λ_G.

From now on, we label the vertices in G by $[n]$, and the u-th row of A_G corresponds to the vertex labeled by u. Recall that the rotation map Rot_G is an exchanging permutation. Let P be the corresponding permutation matrix of size $\mathrm{vol}(G) \times \mathrm{vol}(G)$. That is, for $u, v \in [n]$, $i \in [d_u]$ and $j \in [d_v]$, the $((u, i), (v, j))$-th component of P is 1 if and only if $\mathrm{Rot}_G(u, i) = (v, j)$. We know that $P^2 = I$. For each $u \in [n]$, let Z_u denote the normalized adjacency matrix of graph H_u. Define the following block diagonal matrix

$$Z = \begin{pmatrix} Z_1 & & & \\ & Z_2 & & \\ & & \ddots & \\ & & & Z_n \end{pmatrix}.$$

Each block Z_u has size $[d_u] \times [d_u]$. M can be decomposed to three parts: $M = ZPZ$. Since $\lambda_{\widehat{G}} = 1 - \max_{\alpha \perp \vec{1}} \frac{<\alpha, M\alpha>}{<\alpha,\alpha>}$, then we only have to give an upper bound for $\frac{<\alpha, M\alpha>}{<\alpha,\alpha>}$ assuming $\alpha \perp \vec{1}$.

Note that α has length $\mathrm{vol}(G)$. We divide α into n segments $\alpha_1 \circ \alpha_2 \circ \cdots \circ \alpha_n$, called *segment expression*, where $\circ$ refers to vector concatenation and α_u has length d_u. For each α_u, it can be uniquely decomposed to $\alpha_u^{\|} + \alpha_u^{\perp}$, where $\alpha_u^{\|} = a_u \cdot \vec{1}$ for some real number a_u and $\alpha_u^{\perp} \perp \vec{1}$. [3] We decompose $\alpha = \alpha^{\|} + \alpha^{\perp}$ by letting $\alpha^{\|} = \alpha_1^{\|} \circ \alpha_2^{\|} \circ \cdots \circ \alpha_n^{\|}$ and $\alpha^{\perp} = \alpha_1^{\perp} \circ \alpha_2^{\perp} \circ \cdots \circ \alpha_n^{\perp}$. Since both α and $\alpha^{\perp}$ are orthogonal to $\vec{1}$, $\alpha^{\|} \perp \vec{1}$ either. It means that $< \alpha^{\|}, \vec{1} >= \sum_{u \in [n]} a_u d_u = 0$. In addition, $Z\alpha^{\|} = \alpha^{\|}$, then we have

$$
\begin{aligned}
\frac{< \alpha, M\alpha >}{< \alpha, \alpha >} &= \frac{< (\alpha^{\|} + \alpha^{\perp}), ZPZ(\alpha^{\|} + \alpha^{\perp}) >}{< \alpha, \alpha >} \\
&= \frac{< \alpha^{\|} + Z\alpha^{\perp}, P(\alpha^{\|} + Z\alpha^{\perp}) >}{< \alpha, \alpha >}.
\end{aligned}
$$

Since P is symmetric and $P^2 = I$, P has eigenvalues 1 and -1. Suppose that the multiplicity of 1 is k and that of -1 is $N - k$, where $N = \mathrm{vol}(G)$. Let $\xi_1, \ldots, \xi_N$ be all the orthogonal unit eigenvectors of P, where $\xi_1, \ldots, \xi_k$ are in the 1-eigenspace and $\xi_{k+1}, \ldots, \xi_N$ are in the (-1)-eigenspace. For any $\gamma \in \mathbb{R}^N$, suppose that $\eta = (c_1\xi_1 + \cdots + c_k\xi_k) + (c_{k+1}\xi_{k+1} + \cdots + c_N\xi_N)$. We have $P\gamma = (c_1\xi_1 + \cdots + c_k\xi_k) - (c_{k+1}\xi_{k+1} + \cdots + c_N\xi_N)$. Let $S = \mathrm{SPAN}\{\xi_1, \ldots, \xi_k\}$, then P is a reflection through S. Denote by θ the angle between γ and S. We have $< \gamma, P\gamma >= (\cos 2\theta) \cdot < \gamma, \gamma >$.

By this observation, let θ be the angle between $\alpha^{\|} + Z\alpha^{\perp}$ and S, ϕ be the angle between $\alpha^{\|}$ and S. We also define ψ_1 to be the angle between $\alpha^{\|}$ and $\alpha^{\|} + Z\alpha^{\perp}$, and ψ_2 to be the angle between $\alpha^{\|}$ and $\alpha^{\|} + \alpha^{\perp}$. Noting that $\alpha^{\|}$ is orthogonal to $\alpha^{\perp}$ and $Z\alpha^{\perp}$, the angles $\psi_1, \psi_2 \in [0, \frac{\pi}{2}]$. We choose proper direction of S such that $\phi \in [0, \frac{\pi}{2}]$ either. Clearly, $\theta \in [\phi - \psi_1, \phi + \psi_1]$. We have

$$
\begin{aligned}
\frac{< \alpha, M\alpha >}{< \alpha, \alpha >} &= \cos 2\theta \cdot \frac{< \alpha^{\|} + Z\alpha^{\perp}, \alpha^{\|} + Z\alpha^{\perp} >}{< \alpha^{\|} + \alpha^{\perp}, \alpha^{\|} + \alpha^{\perp} >} \\
&= \cos 2\theta \cdot \frac{\| \alpha^{\|} \|^2}{\| \alpha^{\|} + \alpha^{\perp} \|^2} \cdot \frac{\| \alpha^{\|} + Z\alpha^{\perp} \|^2}{\| \alpha^{\|} \|^2} \\
&= \cos 2\theta \cdot \frac{\cos^2 \psi_2}{\cos^2 \psi_1}.
\end{aligned}
$$

[3] For the simplicity of symbols, we always use $\vec{1}$ to denote an all 1 vector without its length, which is implicated by context.

Then we show the following two sublemmas.

Sublemma 1.

$$\frac{\tan\psi_1}{\tan\psi_2} \le \eta,$$

where $\eta = \max_{H\in\mathcal{H}}\{\eta_H\}$.

Proof.

$$\begin{aligned}
\frac{\tan\psi_1}{\tan\psi_2} &= \frac{\| Z\alpha^{\perp} \|}{\| \alpha^{\perp} \|} \\
&= \frac{1}{\| \alpha^{\perp} \|}\cdot\sqrt{< \alpha^{\perp}, Z^2\alpha^{\perp} >} \\
&= \frac{1}{\| \alpha^{\perp} \|}\cdot\Big(\sum_{u\in[n]} < \alpha_u^{\perp}, Z_u^2\alpha_u^{\perp} >\Big)^{\frac{1}{2}} \\
&\le \frac{1}{\| \alpha^{\perp} \|}\cdot\Big(\sum_{u\in[n]} \eta_u^2 < \alpha_u^{\perp}, \alpha_u^{\perp} >\Big)^{\frac{1}{2}} \\
&\le \max_{H\in\mathcal{H}}\{\eta_H\}\cdot\frac{1}{\| \alpha^{\perp} \|}\cdot\Big(\sum_{u\in[n]} < \alpha_u^{\perp}, \alpha_u^{\perp} >\Big)^{\frac{1}{2}} \\
&= \eta\cdot\frac{1}{\| \alpha^{\perp} \|}\cdot \| \alpha^{\perp} \| \\
&= \eta.
\end{aligned}$$

□

Sublemma 2.

$$\cos 2\phi \le 1 - \lambda_G.$$

Proof. Firstly, we define a T-operation: $\mathbb{R}^N \to \mathbb{R}^n$ as follows. For any $\beta \in \mathbb{R}^N$, let $\beta = \beta_1 \circ \beta_2 \circ \cdots \circ \beta_n$ be the segment expression. Then $T(\beta)$ is defined as $(b_1, \ldots, b_n) \in \mathbb{R}^n$, in which $b_u = \sum_{j\in[d_u]} \beta_{uj}$, where β_{uj} is the j-th component of β_u. We also define the T^{-1}-operation: $\mathbb{R}^n \to \mathbb{R}^N$ as follows. For any $\vec{b} = (b_1, \ldots, b_n) \in \mathbb{R}^n$, $T^{-1}(\vec{b}) = \beta_1 \circ \beta_2 \circ \cdots \circ \beta_n$, in which each component of β_u is $\frac{b_u}{d_u}$.

Recall that $\alpha^{\|} = \alpha_1^{\|} \circ \alpha_2^{\|} \circ \cdots \circ \alpha_n^{\|}$ and each $\alpha_u^{\|} = a_u \cdot \vec{1}$. Let $e_u \in \mathbb{R}^n$ be the unit vector whose u-th coordinate is 1, and 0 for others. Then we have

$$\begin{aligned}\cos 2\phi &= \frac{<\alpha^{\|}, P\alpha^{\|}>}{<\alpha^{\|}, \alpha^{\|}>} \\ &= \frac{1}{\|\alpha^{\|}\|^2} \cdot <\sum_{u\in[n]} a_u d_u \cdot T^{-1}(e_u), P\sum_{u\in[n]} a_u d_u \cdot T^{-1}(e_u)> \\ &= \frac{1}{\|\alpha^{\|}\|^2} \cdot <(a_1,\ldots,a_n), T(P\sum_{u\in[n]} a_u d_u \cdot T^{-1}(e_u))> \\ &= \frac{1}{\|\alpha^{\|}\|^2} \cdot <(a_1,\ldots,a_n), (a_1d_1,\ldots,a_nd_n)D^{-1}A_G> .\end{aligned}$$

The last equality follows from the probability transition among clouds in $\widehat{G}$. Since $\|\alpha^{\|}\|^2 = \sum_{u\in[n]} a_u^2 d_u$, we have

$$\begin{aligned}\cos 2\phi &= \frac{1}{\sum\limits_{u\in[n]} a_u^2 d_u} \cdot <(a_1,\ldots,a_n), (a_1,\ldots,a_n)A_G> \\ &= \frac{\sum\limits_{u\sim v} 2a_u a_v}{\sum\limits_{u\in[n]} a_u^2 d_u}.\end{aligned}$$

Comparing this with the Equation (1) in Section 2.2, and noting that $\sum_{u\in[n]} a_u d_u = 0$, we know that $\cos 2\phi \le 1 - \lambda_G$. $\square$

Then we turn to bound $\cos 2\theta \cdot \frac{\cos^2\psi_2}{\cos^2\psi_1}$. We consider the following two cases.

Case 1: $\psi_1 \le \phi$.

In this case, $0 \le 2(\phi - \psi_1) \le 2\theta \le 2(\phi + \psi_1) \le 2\pi$. So we have $\cos 2\theta \le \max\{\cos 2(\phi-\psi_1), \cos 2(\phi+\psi_1)\}$. Since $(2\pi - 2(\phi+\psi_1)) - 2(\phi-\psi_1) = 2\pi - 4\phi \ge 0$, we know that $\cos 2(\phi-\psi_1) \ge \cos 2(\phi+\psi_1)$ and $\cos 2\theta \le \cos 2(\phi-\psi_1)$. The equality holds if and only if the projections of $\alpha^{\|}$, $\alpha^{\perp}$ and $Z\alpha^{\perp}$ on S are colinear. But this does not always hold. Let $t_1 = \cos 2\phi$ and $t_2 = \frac{\tan\psi_1}{\tan\psi_2}$. Using some trigonometric manipulations and noting that $\frac{1}{\cos^2\psi_1} = \frac{\cos^2\psi_2 + t_2^2\sin^2\psi_2}{\cos^2\psi_2}$ and

$\sin 2\psi_1 = \frac{2\tan\psi_1}{1+\tan^2\psi_1}$, we have

$$\begin{aligned}
\cos 2\theta \cdot \frac{\cos^2\psi_2}{\cos^2\psi_1} &\le \cos 2(\phi - \psi_1) \cdot \frac{\cos^2\psi_2}{\cos^2\psi_1} \\
&= \frac{1}{2}t_1(1-t_2^2) + \frac{1}{2}(1+t_2^2)\cos 2\phi \cos 2\psi_2 + t_2 \sin 2\phi \sin 2\psi_2 \\
&\le \frac{1}{2}t_1(1-t_2^2) + \frac{1}{2}\sqrt{t_1^2(1+t_2^2)^2 + 4(1-t_1^2)t_2^2} \\
&= \frac{1}{2}t_1(1-t_2^2) + \frac{1}{2}\sqrt{t_1^2(1-t_2^2)^2 + 4t_2^2}.
\end{aligned}$$

Let $f(t_1, t_2)$ denote the above function. It is easy to verify that both $\frac{\partial f(t_1,t_2)}{\partial t_1}$ and $\frac{\partial f(t_1,t_2)}{\partial t_2}$ are non-negative. f is an increasing function of t_1 and t_2.

Case 2: $\psi_1 > \phi$.

In this case, $0 \le 2\phi < 2\psi_1 \le \pi$. Then $\cos 2\psi_1 < \cos 2\phi = t_1$. Since $\cos 2\psi_1 = 2\cos^2\psi_1 - 1 = \frac{2\cos^2\psi_2}{t_2^2+(1-t_2^2)\cos^2\psi_2} - 1$, we have

$$\cos^2\psi_2 < \frac{(1+t_1)t_2^2}{(1+t_2^2) - t_1(1-t_2^2)}.$$

So

$$\begin{aligned}
\cos 2\theta \cdot \frac{\cos^2\psi_2}{\cos^2\psi_1} &\le \frac{\cos^2\psi_2}{\cos^2\psi_1} \\
&= t_2^2 + (1-t_2^2)\cos^2\psi_2 \\
&< \frac{2t_2^2}{(1+t_2^2) - t_1(1-t_2^2)}.
\end{aligned}$$

Let $g(t_1, t_2)$ denote the above function.

It can be verified that for any $-1 \le t_1 \le 1$ and any t_2, $f(t_1, t_2) - g(t_1, t_2) \ge 0$. Combining Case 1 and 2, we know that for any $\alpha \perp \vec{1}$, $\frac{<\alpha, M\alpha>}{<\alpha,\alpha>} \le f(t_1, t_2)$. Since

f is an increasing function of t_1 and t_2, by Sublemma 1 and 2,

$$\begin{aligned}
\frac{<\alpha, M\alpha>}{<\alpha,\alpha>} &\leq f(1-\lambda_G,\eta) \\
&\leq f(1-\frac{\Phi^2}{2},\eta) \\
&= \frac{1}{2}(1-\eta^2)(1-\frac{\Phi^2}{2}) + \frac{1}{2}\sqrt{(1-\eta^2)^2(1-\frac{\Phi^2}{2})^2+4\eta^2} \\
&= \frac{1}{2}(1-\eta^2) - \frac{1}{2}(1-\eta^2)\frac{\Phi^2}{2} + \frac{1+\eta^2}{2}\sqrt{1-\frac{(1-\eta^2)^2}{(1+\eta^2)^2}\Phi^2+\frac{(1-\eta^2)^2}{4(1+\eta^2)^2}\Phi^4} \\
&\leq \frac{1}{2}(1-\eta^2) - \frac{1}{2}(1-\eta^2)\frac{\Phi^2}{2} + \frac{1+\eta^2}{2} \\
&= 1-\frac{1}{4}(1-\eta^2)\Phi^2.
\end{aligned}$$

By Equation (2) in Section 2.2, $\lambda_{\widehat{G}} = 1 - \max_{\alpha\perp\vec{1}} \frac{<\alpha,M\alpha>}{<\alpha,\alpha>} \geq \frac{1}{4}(1-\eta^2)\Phi^2$. By the Cheeger inequality, $\mathrm{cond}(\widehat{G}) \geq \frac{\lambda_{\widehat{G}}}{2} \geq \frac{1}{8}(1-\eta^2)\Phi^2$. □

5 Testing conductance of G

In this section, we give a tester for $\mathrm{cond}(G)$ and prove our main theorem.

5.1 Description of our testing algorithm

As we stated in the introduction, we essentially test $\mathrm{cond}(\widehat{G})$ instead. We invoke the tester for the bounded degree model in [5]. Consider the following random walk in the d^2-regular graph $\widehat{G}$: starting from any vertex, in each step, choose an outgoing edge with probability $\frac{1}{2d^2}$, and with the remaining probability $\frac{1}{2}$, it stays at the current vertex. Denote by $\widetilde{G}$ the graph based on $\widehat{G}$ in which each vertex has d^2 more self-loops. Then the random walk defined above is equivalent to the standard random walk on $\widetilde{G}$, and for any $S \subseteq \widehat{V}$, $\mathrm{cond}_{\widehat{G}}(S) = 2\mathrm{cond}_{\widetilde{G}}(S)$. By the definition of the non-uniform Zig-Zag product, a step of random walk from (u, a) crossing edge (i, j) is decomposed to three steps. The "Zig" and "Zag" steps are steps in H_u and H_v, respectively. They do not need to query graph G. The only information we want to query from G is $\mathrm{Rot}_G(u, a')$. So each step of random walk in $\widehat{G}$ requires querying the function Rot_G at most once. At the beginning, the random walk requires choosing uniformly a random node in $\widehat{G}$. Since a node in $\widehat{G}$ corresponds to an endpoint of an edge in G, we randomly choose an edge of

G and randomly choose an endpoint of this edge. Equivalently, we only have to randomly choose a component in the storage list defined in Section 2.1, denoted by (v, b), $b \in [d_v]$, which means that the random walk in $\widehat{G}$ starts from the vertex (v, b). Since each edge $\{u, v\}$ in G appears twice in the list (i.e., (u, a), (v, b) for some $a \in [d_u]$, $b \in [d_v]$), such sample is uniform.

Let Φ, ε and σ be the parameters given in Theorem 1. Let $n = |V|$, $m = |E|$ and $N = 2m$. Let $\mathcal{H}$ be a d-regular Zig-Zag sequence. $\mathcal{H}$ is chosen to be an expander family, where $\eta = \max_{H\in\mathcal{H}}\{\eta_H\}$ is a constant in $(0, 1)$. It is well known that such sequence has been explicitly constructed [9, 10]. Firstly, we state a tester defined on a single vertex.

> Vertex Tester
>
> Given $u \in \widehat{V}$, let $l = \frac{8\ln N}{(1-\eta^2)\Phi^2}$ and $k = 8N^{(1+\sigma)/2}$.
>
> (1) Perform k random walks of length l from u.
>
> (2) Let Q be the number of pairwise collisions of endpoints of the k walks and $K = \binom{k}{2}$.
>
> (3) If $\frac{Q}{K} \geq \frac{1+2N^{-\sigma/4}}{N}$, then output REJECT. Otherwise, output ACCEPT.

Then we state a tester for conductance.

> Conductance Tester
>
> Given $G = (V, E)$, let $k_1 = \Omega(\frac{1}{\varepsilon})$, $k_2 = \Omega(\log\frac{1}{\varepsilon})$ and $\widehat{G} = G \textcircled{z} \mathcal{H} = (\widehat{V}, \widehat{E})$.
>
> (1) Choose k_1 vertices uniformly from $\widehat{V}$.
>
> (2) For each chosen vertex u, run the Vertex Tester for k_2 trials. If at least half of them output REJECT, then output REJECT and halt.
>
> (3) If no vertex in step (2) causes REJECT, then output ACCEPT.

5.2 Proof of our main theorem

Note that $\log m = \mathcal{O}(\log n)$. The query complexity of our tester is $k_1 \cdot k_2 \cdot k \cdot l = \mathcal{O}(\frac{m^{(1+\sigma)/2} \log n \log \frac{1}{\varepsilon}}{\varepsilon \Phi^2})$ and also the running time is.

Then we turn to prove the completeness and soundness for our tester. The technique we use here is summarized as follows. We classify the vertices in $\widehat{G}$ as *strong* and *weak* vertices. The random walks starting from the weak vertices mix very slowly, and collide with each other with high probability. If the rejection probability is small, then the size of weak vertex set in $\widehat{G}$ is small either (otherwise, they can be tested by $k_1 = \Omega(\frac{1}{\varepsilon})$ samples from $\widehat{G}$ easily). On the other hand, a vertex set with large volume in G which determines a cut with small conductance implies a large amount of weak vertices in $\widehat{G}$. So the volume of this set must be small. Then we give a patch-up algorithm dealing with these vertices by modifying at most εm edges in G, such that the resulting graph G' has large conductance.

Let $\rho(u)$ denote the collision probability of two random walks of length l in $\widehat{G}$ starting both from u. Let $\vec{p}_u$ be the probability distribution of the random walk from u at the l-th step. Then $\rho(u) = \sum_v (\vec{p}_u(v))^2$. Let $\Delta_l(u)$ denote the distance of $\vec{p}_u$ from the stationary distribution defined as follows.

$$\Delta_l(u)^2 \doteq \| \vec{p}_u - \frac{\vec{1}}{N} \|^2 = \sum_{v \in \widehat{V}} (\vec{p}_u(v) - \frac{1}{N})^2 = \rho(u) - \frac{1}{N}.$$

We have the following sufficient conditions for Conductance Tester in each round of step (2).

Lemma 4. *(Distinguishing Lemma, [5] Corollary 3.2) The following holds with probability at least* $5/6$. *For any* $u \in \widehat{V}$, *if* $\rho(u) < (1 + N^{-\sigma/4})/N$, *then the majority of the* k_2 *trials of the Vertex Tester running on* u *return ACCEPT. If* $\rho(u) > (1 + 6N^{-\sigma/4})/N$, *then the majority of the* k_2 *trials of the Vertex Tester running on* u *return REJECT.*

The proof of the above lemma is based on the fact that the estimate of $\rho(u)$, which is Q/K, concentrates around $\rho(u)$ with very high probability (independent of the values of l and k, see Lemma 1 of [3]). By the Chernoff bound, when $k_2 = \Omega(\log \frac{1}{\varepsilon})$, a gap of size $5N^{-\sigma/4}/N$ on $\rho(u)$ is large enough to distinguish between the two cases that, for every $v \in \widehat{V}$, majority of the k_2 trials on v return ACCEPT or REJECT, respectively. Using these sufficient conditions, we show the correctness of our tester.

For completeness, we have

Lemma 5. *If* $\mathrm{cond}(G) \geq \Phi$, *then the Conductance Tester accepts* G *with probability at least* $2/3$.

Proof. Since $\mathrm{cond}(G) \geq \Phi$, by Lemma 3, $\lambda_{\widehat{G}} \geq \frac{1}{4}(1-\eta^2)\Phi^2$. Let R be the transition probability matrix of the random walk and $\lambda_{\widetilde{G}}$ be the second eigenvalue of $I - R$, the Laplacian of $\widetilde{G}$. Then $R = \frac{I+M}{2}$ (recall that M is the normalized adjacency matrix of $\widehat{G}$) and $\lambda_{\widetilde{G}} = 1 - \frac{1+(1-\lambda_{\widehat{G}})}{2} = \frac{\lambda_{\widehat{G}}}{2} \geq \frac{1}{8}(1-\eta^2)\Phi^2$. Note that all the eigenvalues of $I - R$ are in $[0, 1]$. We have, for any $u \in \widehat{V}$,

$$
\begin{aligned}
\Delta_l(u)^2 &= \| \vec{p}_u - \frac{\vec{1}}{N} \|^2 \\
&= \| (e_u - \frac{\vec{1}}{N})^\top R^l \|^2 \\
&\leq \| e_u - \frac{\vec{1}}{N} \|^2 \cdot (1-\lambda_{\widehat{G}})^{2l} \\
&\leq (1 - \frac{1}{8}(1-\eta^2)\Phi^2)^{\frac{16 \ln N}{(1-\eta^2)\Phi^2}} \\
&\leq \frac{1}{N^2}.
\end{aligned}
$$

Then $\rho(u) = \Delta_l(u)^2 + \frac{1}{N} < \frac{1+N^{-\sigma/4}}{N}$. By the distinguishing lemma, the tester accepts G with probability at least $2/3$. □

For soundness, we show that, if the tester rejects with probability less than $2/3$, then there exists a graph G' that is ε-close to G and $\mathrm{cond}(G') = \Omega(\Phi^2)$. The following two lemmas for bounded degree graphs are given in [5], and we prove them as follows for self-containment.

Lemma 6. *Let $S \subseteq \widehat{V}$ of size $s \leq \frac{N}{2}$ and $\mathrm{cond}_{\widetilde{G}}(S) \leq \phi$. Then for any $l > 0$, there exists a node $u \in S$ satisfying $\Delta_l(u) \geq \frac{(1-4\phi)^l}{2\sqrt{s}}$.*

Proof. We only have to show that, for a randomly chosen $u \in S$, the expectation of $\Delta_l(u)$, which is $\frac{1}{s}\sum_{u\in S}\Delta_l(u)$, is at least $\frac{(1-4\phi)^l}{2\sqrt{s}}$. Since

$$
\begin{aligned}
\frac{1}{s}\sum_{u\in S}\Delta_l(u) &= \frac{1}{s}\sum_{u\in S} \| \vec{p}_u - \frac{\vec{1}}{N} \| \\
&= \frac{1}{s}\sum_{u\in S} \| (e_u - \frac{\vec{1}}{N})^\top R^l \| \\
&\geq \| (\frac{1}{s}\sum_{u\in S} e_u - \frac{\vec{1}}{N})^\top R^l \| \\
&= \| \alpha^\top R^l \|,
\end{aligned}
$$

where $\alpha = \frac{1}{s}\sum_{u\in S} e_u - \frac{\vec{1}}{N}$, we only have to show that $\| \alpha^\top R^l \|^2 \geq \frac{(1-4\phi)^{2l}}{4s}$.

Let $0 = \lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_N \leq 1$ be all the eigenvalues of $I - R$, and $f_1, f_2, \ldots, f_N$ be the corresponding eigenvectors. Consider

$$\alpha^\top R\alpha = (\frac{1}{s}\sum_{u\in S} e_u - \frac{\vec{1}}{N})^\top R(\frac{1}{s}\sum_{u\in S} e_u - \frac{\vec{1}}{N}).$$

Since

$$\begin{aligned}(\frac{1}{s}\sum_{u\in S} e_u)^\top R(\frac{1}{s}\sum_{u\in S} e_u) &= \frac{1}{s^2}\sum_{u\in S}\sum_{v\in S} e_u^\top R e_v \\ &= \frac{1}{s^2}\cdot\frac{1}{2d^2}\cdot(1-\mathrm{cond}_{\widetilde{\mathrm{G}}}(\mathrm{S}))\mathrm{vol}_{\widetilde{\mathrm{G}}}(\mathrm{S}) \\ &= \frac{1}{s}(1-\mathrm{cond}_{\widetilde{\mathrm{G}}}(\mathrm{S})) \\ &\geq \frac{1}{s}(1-\phi),\end{aligned}$$

and both $(\frac{1}{s}\sum_{u\in S} e_u)^\top R\frac{\vec{1}}{N}$ and $(\frac{\vec{1}}{N})^\top R\frac{\vec{1}}{N}$ equals $\frac{1}{N}$, we have

$$\alpha^\top R\alpha \geq \frac{1}{s}(1-\phi) - \frac{1}{N}.$$

On the other hand, let $\alpha = \sum_{i=1}^N \alpha_i f_i$. Then $\alpha^\top R\alpha = \sum_{i=1}^N \alpha_i^2\lambda_i$. Thus,

$$\sum_{i=1}^N \alpha_i^2\lambda_i \geq \frac{1}{s}(1-\phi) - \frac{1}{N}.$$

Let $B_1 = \{i|\lambda_i \geq 1-4\phi\}$ and $B_2 = [N]\setminus B_1$. We have

$$\sum_{i\in B_1}\alpha_i^2 + \sum_{i\in B_2}\alpha_i^2(1-4\phi) \geq \frac{1}{s}(1-\phi) - \frac{1}{N}.$$

Note that

$$\sum_{i=1}^N \alpha_i^2 = \| \alpha \|^2 = \| \frac{1}{s}\sum_{u\in S} e_u - \frac{\vec{1}}{N} \|^2 = \frac{1}{2} - 2\cdot\frac{1}{N} + \frac{1}{N} = \frac{1}{s} - \frac{1}{N}.$$

We have

$$\sum_{i\in B_1}\alpha_i^2 + (\frac{1}{s} - \frac{1}{N} - \sum_{i\in B_1}\alpha_i^2)(1-4\phi) \geq \frac{1}{s}(1-\phi) - \frac{1}{N}.$$

Solving this, we have

$$\sum_{i \in B_1} \alpha_i^2 \geq \frac{3}{4s} - \frac{1}{N} \geq \frac{1}{4s}.$$

So

$$\| \alpha^\top R^l \|^2 = \alpha^\top R^{2l} \alpha = \sum_{i=1}^{N} \alpha_i^2 \lambda_i^{2l} \geq (1-4\phi)^{2l} \sum_{i \in B_1} \alpha_i^2 \geq \frac{(1-4\phi)^{2l}}{4s}.$$

The lemma follows. $\square$

Lemma 7. *Let $T \subseteq S \subseteq \widehat{V}$, $s = |S| \leq \frac{N}{2}$ and $\mathrm{cond}_{\widehat{G}}(S) \leq \phi$. Let $t = |T| = (1-\delta)s$ for some $\delta \in [0, \frac{1}{5}]$. Then there exists a node $v \in T$ satisfying $\Delta_l(v) \geq (\frac{1}{2} - \sqrt{\frac{\delta}{1-\delta}}) \cdot \frac{(1-4\phi)^l}{\sqrt{s}}$.*

Proof. Let $\alpha = \frac{1}{s} \sum_{u \in S} e_u - \frac{\vec{1}}{N}$ and $\beta = \frac{1}{t} \sum_{v \in T} e_v - \frac{\vec{1}}{N}$. Let $\alpha = \sum_{i=1}^{N} \alpha_i f_i$ and $\beta = \sum_{i=1}^{N} \beta_i f_i$ be the representation of α and β in the basis $\{f_i\}_{i \in [N]}$. Then

$$\| \alpha - \beta \|^2 = (\frac{1}{t} - \frac{1}{s})^2 \cdot t + \frac{1}{s^2} \cdot (s-t) = \frac{1}{t} - \frac{1}{s} = \frac{\delta}{1-\delta} \cdot \frac{1}{s}.$$

By the triangle inequality,

$$\sum_{i \in B_1} \beta_i^2 \geq (\sqrt{\sum_{i \in B_1} \alpha_i^2} - \sqrt{\sum_{i \in B_1} (\alpha_i - \beta_i)^2})^2.$$

Note that $\sum_{i \in B_1} \alpha_i^2 \geq \frac{1}{4s}$, $\sum_{i \in B_1} (\alpha_i - \beta_i)^2 \leq \sum_{i=1}^{N} (\alpha_i - \beta_i)^2 = \frac{\delta}{1-\delta} \cdot \frac{1}{s}$ and $\frac{\delta}{1-\delta} \cdot \frac{1}{s} \leq \frac{1}{4s}$. We have

$$\sum_{i \in B_1} \beta_i^2 \geq (\frac{1}{2\sqrt{s}} - \sqrt{\frac{\delta}{1-\delta}} \cdot \frac{1}{\sqrt{s}})^2 = (\frac{1}{2} - \sqrt{\frac{\delta}{1-\delta}})^2 \cdot \frac{1}{s}.$$

By the proof of Lemma 6, we have

$$\| \beta^\top R^l \| \geq (\frac{1}{2} - \sqrt{\frac{\delta}{1-\delta}}) \cdot \frac{(1-4\phi)^l}{\sqrt{s}},$$

and there exists a node $v \in T$ satisfying $\Delta_l(v) \geq (\frac{1}{2} - \sqrt{\frac{\delta}{1-\delta}}) \cdot \frac{(1-4\phi)^l}{\sqrt{s}}$. $\square$

Then we turn to show that every cut with small conductance in G is unbalanced in the volumes of the two parts.

Lemma 8. *(Unbalanced Cut Lemma) If the Conductance Tester rejects with probability at most $2/3$, then there is a partition of V in G, denoted by $V = S \cup \overline{S}$, such that the following two properties hold:*

(1) $\mathrm{vol}(S) \leq \frac{1}{5}\varepsilon N$*;*

(2) the conductance of the subgraph induced by $\overline{S}$ is $\Omega(\Phi^2)$.

Proof. We partition V recursively as follows. Initially, set $S = \emptyset$ and $\overline{S} = V$. In each step, if there is a subset $A \subseteq \overline{S}$ of volume at most $\frac{\mathrm{vol}(\overline{S})}{2}$ such that the following two hold:

(i) $\mathrm{vol}(A \cup S) \leq \frac{\mathrm{vol}(G)}{2}$,

(ii) $\frac{E(A, \overline{S} \backslash A)}{\mathrm{vol}(A)}$ is less than $c\Phi^2$ for $c = \frac{(1-\eta^2)\sigma}{100}$,

then move A into S. We recursively do this, until no such A can be found.

The second property is guaranteed by our construction. Then we show that $\mathrm{vol}(S) \leq \frac{1}{5}\varepsilon N$. We say that a vertex (u, a) in $\widehat{G}$ is *weak* if $\rho((u, a)) > (1 + 6N^{-\sigma/4})/N$, otherwise is *strong*. Assume that $\mathrm{vol}(S) > \frac{1}{5}\varepsilon N$. Define $\widehat{S} = \{(u, a) \in \widehat{V} | u \in S, a \in [d_u]\}$, which is the union of the clouds in $\widehat{G}$ corresponding to the vertices in S. Then we have $\frac{1}{5}\varepsilon N < \mathrm{vol}(S) = |\widehat{S}| \leq \frac{N}{2}$, and $\mathrm{cond}_{\widetilde{G}}(\widehat{S}) = \frac{1}{2}\mathrm{cond}_{\widehat{G}}(\widehat{S}) = \mathrm{cond}_G(S) < \frac{1}{2}c\Phi^2$. By Lemma 7, choosing $\delta = \frac{1}{10}$, there exist $\frac{1}{10}$ fraction of nodes in $\widehat{G}$ such that for each of them, denoted by (u, a),

$$\Delta_l((u, a)) \geq \frac{(1 - 2c\Phi^2)^l}{3\sqrt{2N}} > \sqrt{\frac{6N^{-\sigma/4}}{N}},$$

and

$$\rho((u, a)) = \Delta_l((u, a))^2 + \frac{1}{N} > \frac{1 + 6N^{-\sigma/4}}{N}.$$

So there are at least $\frac{1}{10}|\widehat{S}| > \frac{1}{50}\varepsilon N$ weak vertices in $\widehat{G}$. By $k_1 = \Omega(\frac{1}{\varepsilon})$ samples from $\widehat{V}$, at least one of such vertices will chosen with probability $\frac{4}{5}$. By Lemma 4, the tester rejects with probability larger than $\frac{4}{5} \cdot \frac{5}{6} = \frac{2}{3}$, which is a contradiction. Then the lemma follows. $\square$

Now we are ready to propose the patch-up algorithm.

Patch-up Algorithm

(1) Partition G into S and $\overline{S}$ satisfying the two properties in Lemma 8.

(2) Remove all the edges incident to some vertex in S.

(3) For each vertex $u \in S$, repeatedly do the following until the degree of u reaches d_u or $d_u - 1$: choose a vertex $v \in \overline{S}$ with probability $\frac{d_v}{\mathrm{vol}_G(\overline{S})}$. If the current degree of v is less than d_v, then add an edge $\{u, v\}$. Otherwise, if the degree is d_v and there is an edge $\{v, w\}$ such that $w \in \overline{S}$, then remove $\{v, w\}$ and add two edges $\{u, v\}$ and $\{u, w\}$. Otherwise, if no such $\{v, w\}$ exists, then re-sample v in $\overline{S}$ and repeat the above process. Denote by G' the resulting graph.

(4) Output G'.

Denote by G'' the resulting graph after step (2). Since $\mathrm{vol}_G(S) \leq \frac{1}{5}\varepsilon N$, for the volumes of G, G'' and G', we have the following relationships. $\mathrm{vol}(G) - \frac{2}{5}\varepsilon N \leq \mathrm{vol}(G'') \leq \mathrm{vol}(G)$, $\mathrm{vol}(G) - \frac{1}{5}\varepsilon N \leq \mathrm{vol}(G') \leq \mathrm{vol}(G)$ and $\mathrm{vol}(G) = N$. To guarantee the implementation of step (3), we show the following lemma.

Lemma 9. *In each stage of each round in step (3), with probability at least $\frac{1}{2} - \frac{3}{5}\varepsilon$, we can find a vertex v in $\overline{S}$ such that the degree of v is less than d_v or there is an edge $\{v, w\}$ such that $w \in \overline{S}$.*

Proof. In step (3), at each stage of each round, we say that a node v is *unsaturated* if the degree of v is less than d_v. Otherwise, v is *saturated*. If the degree of v is more than $\frac{d_v}{2}$, we say that v is *half-saturated*. Otherwise, v is *half-unsaturated*.

Let $X_1 = \{v \in \overline{S} | \deg_{G''}(v) \leq \frac{d_v}{2}\}$, which is the ensemble of half-unsaturated nodes in $\overline{S}$ before step (3). Let $X_2 = \overline{S} \setminus X_1$. At each stage, we consider the following two cases.

Case 1: $\mathrm{vol}_{G''}(X_1) \geq \frac{1}{2}\mathrm{vol}(G'')$.

Let $Y = \{v \in \overline{S} | v \text{ is unsaturated in } G''\text{, but saturated currently.}\}$. Since for every $v \in X_1$, to saturate v, at least $\frac{d_v}{2}$ edges are needed to be added. Then $\mathrm{vol}_G(Y)$ is at most $2\mathrm{vol}_G(S) \leq \frac{2}{5}\varepsilon N$. So with probability at least

$$\frac{\mathrm{vol}_G(X_1) - \mathrm{vol}_G(Y)}{\mathrm{vol}_G(\overline{S})} \geq \frac{\frac{1}{2}\mathrm{vol}(G'') - \frac{2}{5}\varepsilon N}{N} \geq \frac{1}{2} - \frac{3}{5}\varepsilon,$$

a vertex in $X_1 \setminus Y$ that is unsaturated can be chosen.

Case 2: $\mathrm{vol}_{G''}(X_1) < \frac{1}{2}\mathrm{vol}(G'')$.

Note that $\mathrm{vol}(G'')$ is in fact $\mathrm{vol}_{G''}(\overline{S})$. Then $\mathrm{vol}_{G''}(X_2) > \frac{1}{2}\mathrm{vol}(G'')$. Let $Y = \{v \in \overline{S} | v$ is half-saturated in G'', but no (v, w) such that $w \in \overline{S}$ currently.$\}$. For every $v \in X_2$, to remove all edges incident to v within $\overline{S}$, at least $\frac{d_v}{2}$ edges are needed to be removed. Since at most $\frac{1}{10}\varepsilon N$ edges are removed within $\overline{S}$, $\mathrm{vol}_G(Y)$ is at most $\frac{1}{10}\varepsilon N \cdot 2 = \frac{1}{5}\varepsilon N$. So with probability at least

$$\frac{\mathrm{vol}_G(X_2) - \mathrm{vol}_G(Y)}{\mathrm{vol}_G(\overline{S})} \geq \frac{\frac{1}{2}\mathrm{vol}(G'') - \frac{1}{5}\varepsilon N}{N} \geq \frac{1}{2} - \frac{2}{5}\varepsilon,$$

a vertex in $X_2 \setminus Y$ can be chosen.

Combining these two cases, the lemma follows. □

By the above lemma, we only have to choose a small ε, for example, $\varepsilon \leq 5/12$, to guarantee that step (3) can be implemented with probability at least $1/4$. Let $G' = (V', E')$. By the algorithm, the number of modified edges, including added and removed edges, is upper bounded by $2\mathrm{vol}(S) + \frac{1}{2}\mathrm{vol}(S) \leq \varepsilon m$. Note that $|E'| \leq |E| = m$. The distance between G and G' is at most $\frac{\varepsilon m}{\max\{|E|,|E'|\}} = \varepsilon$. Then we only have to show that $\mathrm{cond}(G')$ is large.

Lemma 10. *(Patch-up Lemma)*

$$\mathrm{cond}(G') = \Omega(\Phi^2).$$

Proof. Consider any vertex set A of volume at most $\mathrm{vol}(G')/2$ in G'. Denote by $r = \mathrm{vol}(A) \leq \mathrm{vol}(G')/2$. From now on, whenever we mention the volume of a vertex set without subscript, we mean the volume of this set in G'. Let $\overline{A} = V \setminus A$ denote the complement of A. Consider the following two cases.

Case 1: $\mathrm{vol}(S \cap A) > 2 \cdot \mathrm{vol}(\overline{S} \cap A)$.

Since $\mathrm{vol}(S \cap A) + \mathrm{vol}(\overline{S} \cap A) = r$, we have $\mathrm{vol}(S \cap A) > 2r/3$. Since every edge incident to $S \cap A$ in G' are incident to $\overline{S}$ either, there are at least $\mathrm{vol}(S \cap A) - \mathrm{vol}(\overline{S} \cap A) > \mathrm{vol}(S \cap A)/2$ edges connecting from $S \cap A$ to $\overline{S} \cap \overline{A}$. We have

$$\mathrm{cond}_{G'}(A) > \frac{1}{2} \cdot \frac{2r}{3} \cdot \frac{1}{r} = \frac{1}{3} = \Omega(\Phi^2).$$

Case 2: $\mathrm{vol}(S \cap A) \leq 2 \cdot \mathrm{vol}(\overline{S} \cap A)$.

In this case, $\mathrm{vol}(\overline{S} \cap A) \geq r/3$. We consider the number of edges between $\overline{S} \cap A$ and $\overline{S} \cap \overline{A}$ in G as follows. By the construction of S in Lemma 8,

$$E(\overline{S} \cap A, \overline{S} \cap \overline{A}) \geq c\Phi^2 \cdot \min\{\mathrm{vol}_G(\overline{S} \cap A), \mathrm{vol}_G(\overline{S} \cap \overline{A})\}.$$

Subcase 1: If $\mathrm{vol}_G(\overline{S} \cap A) \leq \mathrm{vol}_G(\overline{S} \cap \overline{A})$, then we have

$$E(\overline{S} \cap A, \overline{S} \cap \overline{A}) \geq c\Phi^2 \cdot \mathrm{vol}_G(\overline{S} \cap A) \geq c\Phi^2 \cdot \mathrm{vol}(\overline{S} \cap A) \geq \frac{1}{3} cr\Phi^2.$$

Subcase 2: Otherwise, $\mathrm{vol}_G(\overline{S} \cap A) > \mathrm{vol}_G(\overline{S} \cap \overline{A})$. Noting that $\mathrm{vol}(\overline{S} \cap A) \leq \mathrm{vol}(G')/2 \leq \mathrm{vol}(G)/2$, we have

$$\begin{aligned}
\mathrm{vol}(\overline{S} \cap \overline{A}) &\geq \mathrm{vol}_{G''}(\overline{S} \cap \overline{A}) \\
&= \mathrm{vol}(G'') - \mathrm{vol}_{G''}(\overline{S} \cap A) \\
&\geq \mathrm{vol}(G) - \frac{2}{5}\varepsilon N - \mathrm{vol}(\overline{S} \cap A) \\
&\geq \frac{1}{2}\mathrm{vol}(G) - \frac{2}{5}\varepsilon N \\
&= (\frac{1}{2} - \frac{2}{5}\varepsilon)N.
\end{aligned}$$

So, when we choose $\varepsilon \leq 5/6$, we have

$$\begin{aligned}
E(\overline{S} \cap A, \overline{S} \cap \overline{A}) &\geq c\Phi^2 \cdot \mathrm{vol}_G(\overline{S} \cap \overline{A}) \\
&\geq (\frac{1}{2} - \frac{2}{5}\varepsilon) c\Phi^2 N \\
&\geq \frac{1}{3} cr\Phi^2.
\end{aligned}$$

Thus, in any subcase, $E(\overline{S} \cap A, \overline{S} \cap \overline{A}) \geq \frac{1}{3} cr\Phi^2$. Note that for any edge $\{u, v\}$ removed from $(\overline{S} \cap A, \overline{S} \cap \overline{A})$, some edge $\{u, w\}$ or $\{v, w\}$ crossing the cut $(A, \overline{A})$ would be added. So we have

$$\mathrm{cond}_{G'}(A) \geq \frac{1}{3} cr\Phi^2 \cdot \frac{1}{r} = \frac{1}{3} c\Phi^2 = \Omega(\Phi^2).$$

Combining Case 1 and 2, the lemma follows. □

Now the proof of Theorem 1 is completed.

References

[1] F. Chung. Spectual graph theory. American Mathematical Society, ISBN 0-8218-0315-8, 1997.

[2] A. Czumaj and C. Sohler. Testing expansion in bounded degree graphs. Proc. 48th Annual Symposium on Foundations of Computer Science, 570-578, 2007.

[3] O. Goldreich and D. Ron. On testing expansion in bounded-degree graphs. Electronic Colloquium on Computational Complexity, TR00-020, 2000.

[4] O. Goldreich and D. Ron. Property testing in bounded degree graphs. Algorithmica, 32(2):302-343, 2002. Conference version in STOC 1997.

[5] S. Kale and C. Seshadhri. Testing expansion in bounded degree graphs. SIAM Journal on Comput. 40(3): 709-720, 2011. Technical Report: ECCC TR07-076, 2007.

[6] R. Motwani and P. Raghavan. Randomized Algorithms. Cambridge University Press, ISBN 0-521-47465-5, 2000.

[7] A. Nachmias and A. Shapira. Testing the expansion of a graph. Electronic Colloquium on Computational Complexity, TR07-118, 2007.

[8] O. Reingold. Undirected connectivity in log-space. Journal of the ACM 55, 4, Article 17 (Sept. 2008), 24 pages.

[9] O. Reingold, S. Vadhan and A. Wigderson, Entropy waves, the Zig-Zag graph product, and new constant-degree expanders. Annals of Maths, 155, 1(January), 2002.

[10] E. Rozenman and S. Vadan. Derandomized Squaring of graphs. International Workshop on Randomization and Approximation Techniques in Computer Science. LNCS 3624, 436-447, 2005.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de