Economical toric spines via Cheeger's Inequality

Noga Alon[*]   Bo'az Klartag[†]

Abstract

Let $G_\infty = (C_m^d)_\infty$ denote the graph whose set of vertices is $\{1, \ldots, m\}^d$, where two distinct vertices are adjacent iff they are either equal or adjacent in C_m in each coordinate. Let $G_1 = (C_m^d)_1$ denote the graph on the same set of vertices in which two vertices are adjacent iff they are adjacent in one coordinate in C_m and equal in all others. Both graphs can be viewed as graphs of the d-dimensional torus. We prove that one can delete $O(\sqrt{d} m^{d-1})$ vertices of G_1 so that no topologically nontrivial cycles remain. This improves an $O(d^{\log_2(3/2)} m^{d-1})$ estimate of Bollobás, Kindler, Leader and O'Donnell. We also give a short proof of a result implicit in a recent paper of Raz: one can delete an $O(\sqrt{d}/m)$ fraction of the edges of G_∞ so that no topologically nontrivial cycles remain in this graph. Our technique also yields a short proof of a recent result of Kindler, O'Donnell, Rao and Wigderson; there is a subset of the continuous d-dimensional torus of surface area $O(\sqrt{d})$ that intersects all nontrivial cycles. All proofs are based on the same general idea: the consideration of random shifts of a body with small boundary and no- nontrivial cycles, whose existence is proved by applying the isoperimetric inequality of Cheeger or its vertex or edge discrete analogues.

1   Introduction

Let $G_\infty = (C_m^d)_\infty$ denote the d-(AND)-power of the cycle C_m on the vertices $M = \{1, 2, \ldots, m\}$, that is, the graph whose set of vertices is M^d, where two distinct vertices $(i_1, i_2, \ldots, i_d)$ and $(j_1, j_2, \ldots, j_d)$ are adjacent iff for every index s, i_s and j_s are either equal or adjacent in C_m. Similarly, let $G_1 = (C_m^d)_1$ denote the graph on the set of vertices M^d in which two vertices $(i_1, i_2, \ldots, i_d)$ and $(j_1, j_2, \ldots, j_d)$ are adjacent iff they are equal in all coordinates but one, in which they are adjacent in C_m.

Both graphs G_∞ and G_1 can be viewed as graphs of the d-dimensional torus. A cycle in any of them is called *nontrivial* if it wraps around the torus, that is, if its projection along at least

one of the coordinates contains the full cycle C_m. A *spine* (or an *edge-spine*) is a set of edges that intersects every nontrivial cycle. It is easy to see that there is a spine in G_∞ containing a fraction of $O(d/m)$ of the edges. A recent result of Raz [7], motivated by the investigation of parallel repetition of the odd cycle game, can be used to show that there are much smaller spines consisting of only a fraction of $O(\sqrt{d}/m)$ of the edges. Here we prove the following sharper version of this result.

Theorem 1 *There exists an edge-spine of G_∞ containing a fraction of at most $2\mu/(3^d - 1) = O(\sqrt{d}/m)$ of the edges of G_∞, where here $\mu = \sqrt{2 \cdot (3^d - 1) \cdot (3^d - (1 + 2\cos(\pi/m))^d)}$.*

It is not difficult to see that the size of the smallest edge-spine in G_1 is precisely dm^{d-1}. Indeed, the set

$$\{ \ \{(i_1, \ldots, i_{s-1}, 0, i_{s+1}, \ldots i_d), (j_1, \ldots, j_{s-1}, 1, j_{s+1}, \ldots, j_d)\} : \ 1 \leq s \leq d, i_r, j_t \in C_m \ \}$$

forms a spine, and there is no smaller spine as the set of all edges of G_1 can be partitioned into dm^{d-1} pairwise edge disjoint nontrivial cycles .

A *vertex-spine* is a set of vertices that intersects every nontrivial cycle. For vertex spines, the smallest size is known for G_∞ and is not known for G_1. Indeed, improving a result of [8], it is proved in [3] that the size of the smallest vertex spine in G_∞ is $m^d - (m - 1)^d$, that is, the vertex-spine consisting of all vertices in which at least one coordinate is 0 is of minimum size. For G_1 the situation is more complicated. It is easy to see that there is a vertex spine consisting of at most dm^{d-1} vertices. This has been improved in [3], where it is shown that there is a vertex spine of size at most $d^{\log_2(3/2)}m^{d-1} \approx d^{0.6}m^{d-1}$. The following result improves this estimate.

Theorem 2 *There exists a vertex spine of G_1 containing at most $2\pi\sqrt{d}m^{d-1}$ vertices.*

The discrete results above have a continuous analogue studied in [5]. Let $\mathbb{T}^d = \mathbb{R}^d / \mathbb{Z}^d$ be the d-dimensional unit torus. We write Vol_d and Vol_{d-1} for the d-dimensional and $(d-1)$-dimensional Hausdorff measures on the unit torus $\mathbb{T}^d$. A loop is a continuous image of the circle. A loop in $\mathbb{T}^d$ is called contractible if it may be continuously deformed to a single point in $\mathbb{T}^d$. A *spine* in $\mathbb{T}^d$ is a subset $S \subset \mathbb{T}^d$ that intersects any non-contractible loop. Clearly, the set

$$S = \{(x_1, \ldots, x_d) \in [0, 1)^d; \exists i \in \{1, \ldots, d\}, x_i = 0\}$$

is a spine, with $\mathrm{Vol}_{d-1}(S) = d$. In [5] it is shown that we can find a much smaller spine.

Theorem 3 ([5]) *There exists a compact spine $S \subset \mathbb{T}^d$ with $\mathrm{Vol}_{d-1}(S) \leq 2\pi\sqrt{d}$.*

In this paper we give relatively short proofs of the above three theorems. The crucial observation is that in all three cases one can apply either the isoperimetric inequality of Cheeger (see, e.g., [6] for a short proof), or its discrete version for vertex boundary (proved in [1]) or for edge boundary (see, e.g., [2]), to obtain a substructure (an induced subgraph in the discrete case, and a body

2

in the continuous case) containing no nontrivial cycles, whose boundary is small with respect to its volume. The required spine is constructed in all cases by pieces of the boundaries of random shifts of this substructure. The proofs given here, while related to the ones given in [7] and [5], are significantly shorter. More importantly, they supply a clear explanation for the choice of the functions whose level sets provide the required substructures, as these appear naturally as eigenfunctions of the corresponding Laplace operators. Indeed, the proof in [5] also produces a spine by combining pieces of boundaries of random shifts of (several) level sets of an appropriate function, but provides no clear intuition to the choice of this function. It also estimates the boundary using a slightly more complicated argument than the one given here, and thus requires a somewhat tedious computation. The proof in [7] uses a different function, yielding a slightly weaker conclusion. As we briefly remark, our approach can be applied to derive similar results for other examples of graphs and bodies.

2 The discrete case

2.1 Edge spines

We start with a review of the discrete version of Cheeger's inequality with Dirichlet boundary conditions. For completeness, we include its proof.

Theorem 4 *Let $G = (V, E)$ be a graph, where $V = \{1, \ldots, n\}$, let $A = (a_{ij})_{i,j \in V}$ be its adjacency matrix and let $Q = diag(d(i)_{i \in V}) - A$ be its Laplace matrix, where $d(i)$ is the degree of i. Let $U \subset V$ be a set of vertices, and let $x = (x_1, x_2, \ldots, x_n)$ be a vector assigning a value x_i to vertex number i. Assume, further that $x_j = 0$ for all $j \in U$ and that for every $W \subset V - U$, $e(W, V - W) \geq c|W|$, where $e(W, V - W)$ is the number of edges joining a vertex of W with one of its complement. Then $x^t Q x \geq \frac{c^2}{2D} \sum_{i \in V} x_i^2$, where D is the maximum degree of a vertex of G.*

Proof: Without loss of generality assume that $V - U = \{1, 2, \ldots, r\}$ and that $x_1^2 \geq x_2^2 \ldots \geq x_r^2$. Since $x^t Q x = \sum_{ij \in E}(x_i - x_j)^2$, by Cauchy Schwartz:

$$2D\Big(\sum_{i \in V} x_i^2\Big) \cdot x^t Q x \geq \sum_{ij \in E}(x_i + x_j)^2 \sum_{ij \in E}(x_i - x_j)^2 \geq \Big[\sum_{ij \in E, i < j}(x_i^2 - x_j^2)\Big]^2. \tag{1}$$

Replacing each term $x_i^2 - x_j^2$ in the last expression by $(x_i^2 - x_{i+1}^2) + (x_{i+1}^2 - x_{i+2}^2) + \cdots + (x_{j-1}^2 - x_j^2)$, the expression obtained from the sum $S = \sum_{ij \in E, i < j}(x_i^2 - x_j^2)$ contains each term of the form $x_i^2 - x_{i+1}^2$ exactly $e(\{1, 2, \ldots, i\}, \{i + 1, \ldots, n\})$ times, and by assumption this number is at least ci for all $i \leq r$. As $x_i = 0$ for $i > r$ this implies that $S \geq \sum_{i \leq n} ci(x_i^2 - x_{i+1}^2) = c\sum_{i \in V} x_i^2$ (where, by definition, $x_{n+1} = 0$). Plugging in (1) the desired result follows. $\square$

Remark: Note that the proof works even if we only assume that $e(W, V - W) \geq c|W|$ for every W which is a level set of the vector $(x_i^2)_{i \in V}$, that is, for every W consisting of all vertices i with $x_i^2 \geq t$.

Let C_m be, as before, the cycle of length m on the set of vertices $M = \{1, 2, \ldots m\}$ (in this order), and let $G_\infty = (C_m^d)_\infty = (V, E)$ denote its d-(AND)-power. Note that G_∞ is $D = 3^d - 1$ regular.

Lemma 5 *There exists a set W of vertices of G_∞ that contains no nontrivial cycles such that $e(W, V - W) \leq \mu|W|$, where μ is as in Theorem 1, and satisfies $\mu/(3^d - 1) = O(\sqrt{d}/m)$.*

Proof: Let $A = (a_{ij})_{i,j \in M}$ be the adjacency matrix of C_m, and let A' be the matrix obtained from it by replacing the last row and last column by the zero vector. Note that the adjacency matrix of G_∞ is $(I + A)^{\otimes d} - I^{\otimes d}$, where for every matrix B, $B^{\otimes d}$ denotes the tensor product of d copies of B. Note also that if x is a vector of length m and $x_m = 0$, and if $x^{\otimes d}$ is the tensor product of d copies of x, then

$$(x^{\otimes d})^t[(I + A)^{\otimes d} - I^{\otimes d}]x^{\otimes d} = (x^{\otimes d})^t[(I + A')^{\otimes d} - I^{\otimes d}]x^{\otimes d},$$

since these two matrices differ only in entries where the contribution to the quadratic form vanishes, as $x_m = 0$.

A simple computation shows that the vector $x = \sin(\pi j/m)_{j \in M}$ (which satisfies $x_m = 0$) is an eigenvector of A' with eigenvalue $\lambda = 2\cos(\pi/m)$. Therefore, $x^{\otimes d}$ is an eigenvector of $(I + A')^{\otimes d} - I^{\otimes d}$ with eigenvalue $\Lambda = (1 + \lambda)^d - 1 = (1 + 2\cos(\pi/m))^d - 1$. By the above discussion this implies that

$$(x^{\otimes d})^t[(I + A)^{\otimes d} - I^{\otimes d}]x^{\otimes d} = \Lambda||x^{\otimes d}||^2,$$

and as the Laplace matrix of G_∞ is $Q = (3^d - 1)I^{\otimes d} - [(I + A)^{\otimes d} - I^{\otimes d}]$ this implies that

$$\frac{(x^{\otimes d})^t Q x^{\otimes d}}{||x^{\otimes d}||^2} = 3^d - 1 - \Lambda.$$

By Theorem 4 we conclude that there is a subset W of the vertices of G_∞ that contains no vertex with any coordinate being m, so that $e(W, V - W) \leq \mu|W|$ with μ as in Theorem 1. The induced subgraph on W contains no nontrivial cycle since $W \subset (M - \{m\})^d$. $\square$

Proof of Theorem 1: Let W be as in Lemma 5, let $v_1, v_2, \ldots$ be a random sequence of vectors in Z_m^d, and define $W_i = v_i + W = \{v_i + w : w \in W\}$ where addition is taken modulo m in each coordinate. By symmetry, the induced subgraph of G_∞ on each W_i is isomorphic to that on W and hence contains no nontrivial cycle. Obviously, with probability 1 there exists a finite s so that $\cup_{i=1}^s W_i = V$. For each i, let E_i be the set of all edges that connect a vertex of $W_i - \cup_{j<i} W_j$ to a vertex outside W_i. The union of all these sets E_i is clearly a spine, as each cycle that uses no edge of this union is contained in a single set W_i. We claim that the expected value of the random variable $\sum |E_i|$ is at most μm^d. To see this, observe that by Lemma 5, if we choose a random vertex of W and a random edge incident with it, the probability that this edge leads to $V - W$ is at most $\frac{\mu|W|}{(3^d-1)|W|} = \frac{\mu}{3^d-1}$. Fix a vertex $v \in M^d$, and let i be the smallest j so that $v \in W_j$. Conditioning on i being the smallest such j, v is a uniform random vertex of W_i, and hence if we

4

now choose a random edge incident with it, the probability it leads to a vertex outside W_i is at most $\frac{\mu}{3^d-1}$. It follows that the expected size of E_i is at most the expected size of $W_i - \cup_{j<i}W_j$ times μ. Summing over all values of i and using the fact that with probability 1 the union of all sets W_i is V we conclude that the expected value of $\sum |E_i|$ is at most $\mu|V| = \mu m^d$. Thus there is a choice of sets W_i so that the spine $\cup_i E_i$ they provide is of size at most μm^d, completing the proof. $\square$

2.2 Vertex spines

We need the following version of the inequality of [1] with Dirichlet boundary condition. This is an analog of Theorem 4, dealing with vertex boundary instead of edge boundary. Its proof, which is based on the arguments in [1], is somewhat more complicated than that of Theorem 4.

Theorem 6 *Let $G = (V, E)$ be a graph, where $V = \{1, \ldots, n\}$, let $A = (a_{ij})_{i,j \in V}$ be its adjacency matrix and let $Q = diag(d(i)_{i \in V}) - A$ be its Laplace matrix, where $d(i)$ is the degree of i. Let $U \subset V$ be a set of vertices, and let $x = (x_1, x_2, \ldots, x_n)$ be a vector assigning a value x_i to vertex number i. Assume, further that $x_j = 0$ for all $j \in U$ and that for every $W \subset V - U$, $|N(W) - W| \geq c|W|$, where $N(W)$ is the set of all vertices that have a neighbor in W. Then $x^t Q x \geq \frac{c^2}{4+2c^2} \sum_{i \in V} x_i^2$.*

Proof: Put $Y = V - U$. We claim that there is an orientation $\overline{E}$ of E and a function $h : \overline{E} \mapsto [0, 1]$ so that the sum $\sum_{j,(i,j)\in\overline{E}} h(i,j)$ is at most $1 + c$ for all $i \in Y$, the sum $\sum_{j,(j,i)\in\overline{E}} h(j,i)$ is at most 1 for all i and the difference $\sum_{j,(i,j)\in\overline{E}} h(i,j) - \sum_{j,(j,i)\in\overline{E}} h(j,i)$ is at least c for every $i \in Y$. To prove this claim, consider the network flow problem in which the set of vertices consists of a source s, a sink t, a set Y' consisting of a copy y' of every $y \in Y$ and a set V" consisting of a copy v" of every vertex $v \in V$. For each $y' \in Y'$, (s, y') is an arc of the network with capacity $1 + c$, for each v" $\in V$", $(v", t)$ is an arc of the network with capacity 1, and in addition, $(u', u")$ is an arc of capacity 1 for each $u \in Y$, and for each edge uv of G, with $u \in Y$, $v \in V$, the arc $(u', v")$ belongs to the network, and has capacity 1. (Note that if v is also in Y, then the arc $(v', u")$ is also in the network.) It is not difficult to check that the value of the maximum flow in this network is $(1 + c)|Y|$. Indeed, suppose we are given a cut and let $X \subseteq Y$ be the set of all vertices $y \in Y$ such that (s, y') belongs to the cut. Then the cut must contain, for each $v \in (Y - X) \cup N(Y - X)$, at least one arc incident with v". As these arcs are pairwise distinct and there are at least $(1 + c)|Y - X|$ of them, each having capacity 1, it follows that the total capacity of the cut is at least $(1 + c)|X| + (1 + c)|Y - X| = (1 + c)|Y|$. By the Maxflow-Mincut Theorem there exists a flow of value at least $(1 + c)|Y|$, and this is clearly a maximum flow that saturates all edges (s, y') with $y \in Y$. If there is a positive flow in two arcs $(i', j")$ and $(j', i")$ (for some $i, j \in Y$), subtract the minimum of these two from both, to ensure that at least one of these two quantities is zero, and subtract this minimum from the value of the flow on $(s, i'), (j", t)$ and on $(s, j'), (i", t)$, thus keeping it a valid flow without changing the value of the difference between the total flow leaving i' and the total flow going into $i"$. Let h' be the resulting flow. If

5

$h'(i', j") > 0$ for $ij \in E$, orient the edge ij from i to j (in case $h'(i', j") = h'(j', i") = 0$ orient the edge arbitrarily). Finally, for each oriented edge (i, j), define $h(i, j) = h'(i', j")$. One can easily check that the function h satisfies the assertion of the claim.

We next note that the properties of h imply the following two inequalities.

$$\sum_{(i,j)\in\overline{E}} h^2(i, j)(x_i + x_j)^2 \le (4 + 2c^2)\sum_i x_i^2. \tag{2}$$

$$\sum_{(i,j)\in\overline{E}} h(i, j)(x_i^2 - x_j^2) \ge c\sum_i x_i^2. \tag{3}$$

Indeed, (2) follows, as

$$\sum_{(i,j)\in\overline{E}} h^2(i, j)(x_i + x_j)^2 \le 2\sum_{(i,j)\in\overline{E}} h^2(i, j)(x_i^2 + x_j^2)$$

$$= 2\sum_{i\in Y} x_i^2 \Big(\sum_{j,(i,j)\in\overline{E}} h^2(i, j) + \sum_{j,(j,i)\in\overline{E}} h^2(j, i)\Big) \le 2(2 + c^2)\sum_i x_i^2,$$

where here we used the fact that $x_i = 0$ for all $i \notin Y$ and the fact that the sum of squares of reals in $[0, 1]$ whose sum is at most $(1 + c)$ does not exceed $1 + c^2$ (and the sum of squares of real numbers in $[0, 1]$ whose sum is at most 1 does not exceed 1).

To prove (3) note that

$$\sum_{(i,j)\in\overline{E}} h(i, j)(x_i^2 - x_j^2) = \sum_{i\in Y} x_i^2 \Big(\sum_{j,(i,j)\in\overline{E}} h(i, j) - \sum_{j,(j,i)\in\overline{E}} h(j, i)\Big) \ge c\sum_i x_i^2.$$

We can now complete the proof of the theorem. Indeed, by Cauchy-Schwartz, (2) and (3):

$$\frac{x^t Q x}{\sum_i x_i^2} = \frac{\sum_{(i,j)\in\overline{E}}(x_i - x_j)^2}{\sum_i x_i^2} = \frac{\sum_{(i,j)\in\overline{E}}(x_i - x_j)^2 \sum_{(i,j)\in\overline{E}} h^2(i, j)(x_i + x_j)^2}{\sum_i x_i^2 \sum_{(i,j)\in\overline{E}} h^2(i, j)(x_i + x_j)^2}$$

$$\ge \frac{(\sum_{(i,j)\in\overline{E}} h(i, j)(x_i^2 - x_j^2))^2}{(4 + 2c^2)(\sum_i x_i^2)^2} \ge \frac{c^2}{4 + 2c^2},$$

completing the proof. $\square$

Returning to the graph $G_1 = (C_m^d)_1$ defined in the introduction, note that it is a $2d$-regular graph on m^d vertices. If the adjacency matrix of C_m is A, as before, then the one of G_1 is the sum of d terms, each of which is a tensor product of $d - 1$ copies of I and one copy of A. Thus, here, too, we can use the vector $x^{\otimes d}$, where x is as in Lemma 5, and prove the following.

Lemma 7 *There exists a set of vertices W of G_1 that contains no nontrivial cycles so that its vertex boundary in G_1 is of size $c|W|$, and $\frac{c}{1+c} \le 4\sqrt{d}\sin(\frac{\pi}{2m}) \le 2\pi\frac{\sqrt{d}}{m}$.*

6

Proof: The Laplace matrix L of G_1 and the vector $x^{\otimes d}$ above satisfy

$$\frac{(x^{\otimes d})^t L x^{\otimes d}}{||x^{\otimes d}||^2} = 2d - 2d\cos(\pi/m) = 4d\sin^2\frac{\pi}{2m}.$$

By Theorem 6 this implies that there is a set of vertices W containing no vertices with any coordinate being m so that if $|N(W) - W| = c|W|$, then $\frac{c^2}{4+2c^2} \leq 4d\sin^2\frac{\pi}{2m}$. The desired result follows, as

$$\frac{1}{4}\left(\frac{c}{1+c}\right)^2 \leq \frac{c^2}{4+2c^2} \leq 4d\sin^2\frac{\pi}{2m} < \frac{d\pi^2}{m^2},$$

completing the proof. $\square$

Proof of Theorem 2: The proof is very similar to that of Theorem 1. Let W be as in Lemma 7, let $v_1, v_2, \ldots$ be a random sequence of vectors in Z_m^d, and define $W_i = v_i + W = \{v_i + w : w \in W\}$ where addition is taken modulo m in each coordinate. By symmetry, the induced subgraph of G_1 on each W_i is isomorphic to that on W and hence contains no nontrivial cycle. Obviously, with probability 1 there exists a finite s so that the union $\cup_{i=1}^s (W_i \cup N(W_i))$ covers all vertices of G_1. For each i, define $B_i = (N(W_i) - W_i) - \cup_{j<i}(W_j \cup N(W_j))$. The union of all the sets B_i is a vertex spine, as each cycle that uses no vertex of this union is contained in a single set W_i. We claim that for each fixed vertex v of G_1, the probability that v belongs to the above union is $\frac{c}{1+c}$. Indeed, if i is the smallest j so that $v \in (W_j \cup N(W_j))$, v is a uniform random vertex of $W_i \cup N(W_i)$ and the probability that it lies in $N(W_i) - W_i$ is thus precisely $c/(1+c)$, as claimed. By linearity of expectation, the expected size of the union of all sets B_i is $\frac{c}{1+c}m^d$, and the desired result follows. $\square$

Remarks:

- A simple computation shows that for large m and d, the expression $\frac{2\mu}{3^d-1}$ in Theorem 1 is at most $(1 + o(1))\sqrt{\frac{8}{3}}\frac{\pi\sqrt{d}}{m}$.

- By the remark following the proof of Theorem 4, the set W in Lemma 5 is a level set of the vector $x^{\otimes d}$ (or equivalently, the vector obtained from it by squaring each coordinate.)

- Theorem 1 gives an alternative proof of the main result of Raz [7], showing that in the parallel repetition for the maxcut game on an odd cycle, $\Theta(m^2)$ repetitions are required to ensure a value smaller than $1/2$.

- The assertion of Theorems 1 and 2 can be extended to powers of other Cayley graphs.

3   The continuous case

The discrete results above have a continuous analogue. Let d be a dimension, and let $\mathbb{T}^d = \mathbb{R}^d/\mathbb{Z}^d$ be the unit torus. The torus $\mathbb{T}^d$, which will be identified as a set with $[0,1)^d$, inherits

the Riemannian structure of $\mathbb{R}^d$. Recall that we write Vol_d and Vol_{d-1} for the d-dimensional and $(d-1)$-dimensional Hausdorff measures on the unit torus $\mathbb{T}^d$. A *body* here means a non-empty compact set that equals the closure of its interior. A *smooth* function or surface always means here C^∞-smooth.

Lemma 8 *There exists a body $D \subset (0,1)^d \subset \mathbb{T}^d$ with a smooth boundary, such that*

$$Vol_{d-1}(\partial D) \leq 2\pi\sqrt{d}\, Vol_d(D).$$

Proof: Denote

$$h = \inf_{A \subset (0,1)^d} \frac{\mathrm{Vol}_{d-1}(\partial A)}{\mathrm{Vol}_d(A)} \tag{4}$$

where the infimum runs over all bodies A with a smooth boundary in $(0,1)^d$. Cheeger's inequality with Dirichlet boundary conditions on the cube states that for any smooth function $\varphi : [0,1]^d \to \mathbb{R}$ that vanishes on the boundary,

$$-\int_{[0,1]^d} \varphi \triangle \varphi \geq \frac{h^2}{4} \int_{[0,1]^d} \varphi^2, \tag{5}$$

where $\triangle$ is the Laplacian. For a short proof of Cheeger's inequality, see, e.g., [6], Chapter III. The best function to substitute in (5) is the Laplacian eigenfunction $\varphi(x) = \prod_{i=1}^d \sin(\pi x_i)$, which satisfies $\triangle\varphi = -d\pi^2\varphi$. From (5) we thus learn that $d\pi^2 \leq h^2/4$, and the lemma follows. $\square$

Remark. The set $D \subset (0,1)^d$ in Lemma 8 may be chosen to be convex. In fact, as the proof of Cheeger's inequality shows, the set D may be chosen to be a level set of the concave function $\log\varphi(x) = \sum_{i=1}^d \log \sin(\pi x_i)$.

Proof of Theorem 3: Let $v_1, v_2, \ldots \in \mathbb{T}^d$ be a sequence of independent random vectors, uniformly distributed in the torus $\mathbb{T}^d$. Let D stand for the body from Lemma 8, and write $D_i = v_i + D$, where addition is carried in the group $\mathbb{T}^d$. Consider the disjoint union

$$S = \bigcup_{i=1}^\infty S_i, \quad \text{where} \quad S_i = \partial D_i \setminus \cup_{j=1}^{i-1} D_j.$$

Since D has a non-empty interior, then with probability one, $\mathbb{T}^d$ is the union of finitely many D_j's. Thus, if a loop in $\mathbb{T}^d$ does not intersect S, it must be contained in $D_i \subset v_i + (0,1)^d$ for some i, and hence it is contractible. Consequently, S is a compact spine with probability one. It remains to show that $\mathbb{E}\mathrm{Vol}_{d-1}(S) \leq 2\pi\sqrt{d}$. Note that S is a finite union of S_i's, and each S_i is a relatively open subset of the smooth hypersurface ∂D_i. Therefore,

$$\mathrm{Vol}_{d-1}(S) = \lim_{\varepsilon \to 0^+} (2\varepsilon)^{-1}\mathrm{Vol}_d\left(\bigcup_{i=1}^\infty (S_i)_\varepsilon\right) = \lim_{\varepsilon \to 0^+} \varepsilon^{-1}\mathrm{Vol}_d\left(\bigcup_{i=1}^\infty ((S_i)_\varepsilon \cap D_i)\right), \tag{6}$$

where $(S_i)_\varepsilon$ is the set of all points in $\mathbb{T}^d$ whose geodesic distance from S_i is smaller than ε. Fix a point $x \in \mathbb{T}^d$. Then

$$\lim_{\varepsilon \to 0^+} \mathbb{E}\, \varepsilon^{-1}\mathrm{Vol}_d \left(\bigcup_{i=1}^{\infty} ((S_i)_\varepsilon \cap D_i) \right) = \lim_{\varepsilon \to 0^+} \varepsilon^{-1}\mathbb{P}\left(x \in \bigcup_{i=1}^{\infty} ((S_i)_\varepsilon \cap D_i) \right) \qquad (7)$$

There exists a minimal index i such that $x \in D_i$. Let ℓ be this minimal index (so ℓ is a random variable). The crucial observation is that $x - v_\ell$ is distributed uniformly in D. Hence, we may continue (7) with

$$= \lim_{\varepsilon \to 0^+} \varepsilon^{-1}\mathbb{P}\left(x \in \bigcup_{i=\ell}^{\infty} ((S_i)_\varepsilon \cap D_i) \right) \leq \lim_{\varepsilon \to 0^+} \varepsilon^{-1}\mathbb{P}\left(x \in (\partial D_\ell)_\varepsilon \right) = \frac{\mathrm{Vol}_{d-1}(\partial D_\ell)}{\mathrm{Vol}_d(D_\ell)} \leq 2\pi\sqrt{d}, \qquad (8)$$

according to Lemma 8, since $\mathbb{P}(x \in (\partial D_\ell)_\varepsilon) = \mathbb{P}(x - v_\ell \in (\partial D)_\varepsilon)$. From (6) and Fatou's lemma,

$$\mathbb{E}\,\mathrm{Vol}_{d-1}(S) \leq \lim_{\varepsilon \to 0^+} \mathbb{E}\, \varepsilon^{-1}\mathrm{Vol}_d \left(\bigcup_{i=1}^{\infty} ((S_i)_\varepsilon \cap D_i) \right) \leq 2\pi\sqrt{d},$$

where the last inequality follows from (7) and (8). The proof is complete. □

Remarks:

- A spine $S \subset \mathbb{T}^d$ is called regular if it is contained in a finite union of smooth hypersurfaces in $\mathbb{T}^d$. A spine $S \subset \mathbb{T}^d$ is minimal if for any $x \in S$ and $\varepsilon > 0$, the set $S \setminus B(x, \varepsilon)$ is no longer a spine, where $B(x, \varepsilon)$ is the open ball of radius ε about x. By Zorn's lemma, for any compact spine S there exists a minimal sub-spine $S' \subset S$. We may thus assume that the spine S in Theorem 3 is minimal and regular. When S is a minimal regular spine, the set $\mathbb{T}^d \setminus S$ is necessarily connected, and since it intersects all non-contractible loops, it is simply-connected. Note that for a minimal spine S, the set $\bar{S} = \{x \in \mathbb{R}^d; x \mod \mathbb{Z}^d \in S\}$ is the boundary of a $\mathbb{Z}^d$-periodic tiling of $\mathbb{R}^d$ with connected cells of volume one.

- Theorem 3 is tight, up to the value of the constant 2π. Indeed, suppose $S \subset \mathbb{T}^d$ is a minimal spine. Consider the set $\bar{S} = \{x \in \mathbb{R}^d; x \mod \mathbb{Z}^d \in S\}$, and pick a connected component C of $\mathbb{R}^d \setminus \bar{S}$. Then $\mathrm{Vol}_d(C) = 1$ and $\mathrm{Vol}_{d-1}(\partial C) = 2\mathrm{Vol}_{d-1}(S)$. By the classical isoperimetric inequality in $\mathbb{R}^d$,

$$\mathrm{Vol}_{d-1}(S) = \mathrm{Vol}_{d-1}(\partial C)/2 \geq \kappa_d \mathrm{Vol}_d(C)^{(d-1)/d}/2 = \kappa_d/2$$

where $\kappa_d = d\sqrt{\pi}\Gamma(1 + d/2)^{-1/d} \geq \sqrt{d}\left(\sqrt{2\pi e} + o(1)\right)$.

- Our proof uses very few properties of the torus. A straightforward generalization of Theorem 3 might read as follows: Suppose a Lie group G acts transitively by isometries on a simply-connected Riemannian manifold Ω (in our case $G = \Omega = \mathbb{R}^d$). Let Γ be a discrete, co-compact subgroup (in our case $\Gamma = \mathbb{Z}^d$), and let $T \subset \Omega$ be a fundamental domain (in

9

our case $T = [0,1)^d$). Assume that T is simply connected, and write λ for the minimal eigenvalue of minus the Laplacian with Dirichlet boundary conditions on T. Then, there exists a compact spine in Ω/Γ whose surface area is at most $2\sqrt{\lambda}$.

Note that there clearly exists a trivial spine in Ω/Γ whose area is at most $\text{Vol}_{d-1}(\partial T)$. Only in the case where $\sqrt{\lambda} << \text{Vol}_{d-1}(\partial T)$ we obtain a non-trivial conclusion.

- A short argument leading from the continuous Theorem 3 to the discrete Theorem 1 (with a slightly worse constant) appears in [4], Theorem 3.1.

References

[1] N. Alon, Eigenvalues and expanders, Combinatorica 6(1986), 83-96.

[2] N. Alon, O. Schwartz and A. Shapira, An elementary construction of constant-degree expanders, Proc. of the Eighteenth Annual ACM-SIAM SODA (2007), 454-458. Also: Combinatorics, Probability and Computing, to appear.

[3] B. Bollobás, G. Kindler, I. Leader, and R. O'Donnell, Eliminating cycles in the discrete torus, LATIN 2006: Theoretical informatics, 202–210, Lecture Notes in Comput. Sci., 3887, Springer, Berlin, 2006. Also: Algorithmica 50 (2008), no. 4, 446–454.

[4] U. Feige, G. Kindler, R. O'Donnell, Understanding Parallel Repetition Requires Understanding Foams. Proceedings of the Twenty-Second Annual IEEE Conference on Computational Complexity, 2007, 179-192.

[5] G. Kindler, R. O'Donnell, A. Rao and A. Wigderson, Rounding schemes and cubical tilings with sphere-like surface area, Proc. FOCS 2008, to appear.

[6] R. Schoen and S.T. Yau, Lectures on differential geometry, International Press, Cambridge, MA, 1994.

[7] R. Raz, A counterexample to strong parallel repetition, Proc. FOCS 2008, to appear.

[8] M. Saks, A. Samorodnitsky and L. Zosin, A lower bound on the integrality gap for minimum multicut in directed networks, Combinatroica 24 (2004), 525-530.