Eliminating cycles in the discrete torus

Béla Bollobás
Dept. of Mathematical Sciences
University of Memphis
bollobas@msci.memphis.edu

Guy Kindler
Theory Group
Microsoft Research
gkindler@microsoft.com

Imre Leader
Centre for Mathematical Sciences
University of Cambridge
leader@dpmms.cam.ac.uk

Ryan O'Donnell
Theory Group
Microsoft Research
odonnell@microsoft.com

September 22, 2005

Abstract

In this paper we consider the following question: how many vertices of the discrete torus must be deleted so that no topologically nontrivial cycles remain?

We look at two different edge structures for the discrete torus. For $(\mathbb{Z}_m^d)_1$, where two vertices in $\mathbb{Z}_m$ are connected if their L_1 distance is 1, we show a nontrivial upper bound of $d^{\log_2(3/2)} m^{d-1} \approx d^{.6} m^{d-1}$ on the number of vertices that must be deleted. For $(\mathbb{Z}_m^d)_\infty$, where two vertices are connected if their L_∞ distance is 1, Saks, Samorodnitsky and Zosin [8] already gave a nearly tight lower bound of $d(m-1)^{d-1}$ using arguments involving linear algebra. We give a more elementary proof which improves the bound to $m^d - (m-1)^d$, which is precisely tight.

1 Introduction

In this paper we consider a "vertex multicut" problem on discrete torus graphs. Let us begin by defining the two graphs of interest to us.

Definition 1. *The* L_1 discrete torus of width m and dimension d, *denoted* $(\mathbb{Z}_m^d)_1$, *is the undirected graph on vertex set* $\mathbb{Z}_m^d$ *in which two vertices are connected if their* L_1 *distance is 1.*

The L_∞ discrete torus of width m and dimension d, *denoted* $(\mathbb{Z}_m^d)_\infty$, *is the undirected graph on vertex set* $\mathbb{Z}_m^d$ *in which two vertices are connected if their* L_∞ *distance is 1.*

We will also write $(\mathbb{Z}^d)_1$ *and* $(\mathbb{Z}^d)_\infty$ *for the similarly defined infinite graphs on vertex set* $\mathbb{Z}^d$.

In each of these tori we are interested in the set of cycles that "wrap around" the torus in at least one dimension. Let us define this notion formally.

Definition 2. *A cycle in* $(\mathbb{Z}_m^d)_1$ *(respectively,* $(\mathbb{Z}_m^d)_\infty$*) is said to be noncontractible if, when regarded as a loop inside the solid torus, it is homotopically nontrivial.*

Main problem. In this paper we want to study the minimal number of vertices in either discrete torus that must be deleted so that every noncontractible cycle is broken. In other words, we consider the problem of finding the set of vertices of minimal size that intersects every noncontractible cycle in $(\mathbb{Z}_m^d)_1$ or $(\mathbb{Z}_m^d)_\infty$. We denote the minimal number for $(\mathbb{Z}_m^d)_1$ by $S_1(m,d)$ and the minimal number for $(\mathbb{Z}_m^d)_\infty$ by $S_\infty(m,d)$.

We note that there are some obvious bounds that hold for both $S_1(m,d)$ and $S_\infty(m,d)$. A lower bound of m^{d-1} follows by considering, in either graph, the m^{d-1} noncontractible cycles which are parallel to the first axis. These cycles are vertex-disjoint, so at least one vertex must be deleted from each of them. An obvious upper bound of $m^d - (m-1)^d$ is obtained by deleting the union of d "walls", one in each dimension; by a "wall" we mean a set of the form $\{x : x_i = a\}$ for some $i \in [d]$, $a \in \mathbb{Z}_m$.

1.1 History and motivation

The problem discussed in this paper is a natural one in the context of the combinatorics of the discrete torus (see e.g. [2, 1, 3]), but it has other motivations as well.

Discrete foams. Our problem is related to the isoperimetry of periodic tilings of space. The connection is apparent from the following formulation of our problem. We say that a finite set S in $\mathbb{Z}^d$ *generates a discrete foam for* $(\mathbb{Z}^d)_1$ *with periodicity* $m \cdot \mathbb{Z}^d$ if the set

$$\mathbb{Z}^d \setminus \{S + v\}_{v \in m \cdot \mathbb{Z}^d}$$

contains no paths in $(\mathbb{Z}^d)_1$ of infinite length. (We can give a similar definition for $(\mathbb{Z}^d)_\infty$.) It can easily be verified that our problem is identical to that of finding the minimal size of a set generating a discrete foam with periodicity $m \cdot \mathbb{Z}^d$.

This problem can be essentially regarded as that of finding a tiling of $\mathbb{Z}^d$ with periodicity $m \cdot \mathbb{Z}^d$ that has minimal vertex boundary; this is a discrete version of the problem of finding a (continuous) closed foam in $\Re^d$ with periodicity $\mathbb{Z}^d$ and minimal surface area. Although there has been a lot of work on soap bubble and foam problems in $\Re^d$ and even on the flat torus — see e.g. [7] — very little is known. We hope that discrete versions of the problem may prove to be a useful source for new observations regarding foams.

Directed minimum multicut. Another area in which our problem arises is in theoretical computer science, as was noted in a paper of Saks, Samorodnitsky and Zosin [8]. This paper studied the integrality gap of the natural linear programming formulation of the "directed minimum multicut" problem. This is the problem in which one is given a directed graph and d "source-sink" pairs of vertices $(s_1, t_1), \ldots, (s_d, t_d)$, and one is required to delete as few edges as possible so that there is no longer any s_i-to-t_i path. To obtain their integrality gap bound, Saks et al. translated the directed minimum multicut problem on a certain graph to an undirected vertex-deletion problem. Specifically, they looked at the graph $([m]^d)_\infty$ — i.e., the d-dimensional, width m *grid* with L_∞ edges — and studied following quantity:

Definition 3. *$S'_\infty(m,d)$ is the minimum number of vertices in $([m]^d)_\infty$ that need to be deleted to disconnect all d pairs of opposing walls.*

Clearly $S'_\infty(m,d) \geq S_\infty(m,d)$. Saks et al. proved a lower bound of $d(m-1)^{d-1}$ on $S'_\infty(m,d)$, but their proof immediately gives the same lower bound for $S_\infty(m,d)$. This result yielded an integrality gap arbitrarily close to d (which is the best possible) for the directed multicut problem. In this paper we improve the lower bound for $S_\infty(m,d)$ (and thus for $S'_\infty(m,d)$) to $m^d-(m-1)^d$, which exactly matches the upper bound mentioned earlier.

Parallel repetition on odd cycles. Our original motivation came from a problem in the study of parallel repetition of two-prover one-round games [4, 6], and in particular a question due to Feige [5] about how the max-cut problem on odd cycles behaves under parallel repetition.

The details of this problem are beyond the scope of this paper; suffice it to say that it can be reduced to a problem very similar to that of eliminating cycles in $(\mathbb{Z}_m^d)_\infty$ (we give more details in Section 3). However, it seems that solving that problem requires a proof of a lower bound on $S_\infty(m,d)$ that is "robust", in the sense that it should imply a nontrivial bound even under a certain relaxed hypothesis. The lower bound of Saks et al. relies on a linear algebraic argument, and this seems too fragile to give anything once hypotheses are relaxed. Our lower bound, on the other hand, is proven using more elementary methods; hence it seems to have more of a chance to be generalizable.

1.2 Our results

We have two main results. Our first result is an improved upper bound on $S_1(m,d)$.

Theorem 1. $S_1(m,d) \leq d^{\log_2(3/2)} m^{d-1}$.

As far as we know, no nontrivial upper bound on $S_1(m,d)$ was previously known.

Our second result is a lower bound on $S_\infty(m,d)$ that precisely matches the obvious upper bound already discussed. This result improves on the lower bound of Saks, Samorodnitsky and Zosin [8] and eliminates their use of linear algebra.

Theorem 2. $S_\infty(m,d) \geq m^d-(m-1)^d$*, and hence* $S_\infty(m,d) = m^d-(m-1)^d$.

2 The upper bound on $S_1(m,d)$

Our main goal in this section is to prove Theorem 1, showing an upper bound for $S_1(m,d)$. Before doing this, we will motivate our bound by giving a tight construction in two dimensions which has size about $(3/2)m$.

2.1 A tight bound for $(\mathbb{Z}_m^2)_1$

It is easy to see that the following set of size at most $(3/2)m$ blocks all noncontractible cycles in $(\mathbb{Z}_m^2)_1$:

$$S = \{(x,x) : x \in \mathbb{Z}_m\} \cup \{(x,-x) : 0 \leq x \leq k/2\}.$$

Let us sketch a proof of this fact. It is well-known that in two dimensions, $(\mathbb{Z}_m^2)_\infty$ is dual to $(\mathbb{Z}_m^2)_1$. The set S contains a cycle in $(\mathbb{Z}_m^2)_\infty$ that winds once in the first dimension and no times in the second dimension — call such a cycle a $(1,0)$-cycle. This blocks all cycles in $(\mathbb{Z}_m^2)_1$ except those of type $(c,0)$. But S also contains a $(0,1)$-cycle in $(\mathbb{Z}_m^2)_\infty$, thus blocking all $(c,0)$-cycles in $(\mathbb{Z}_m^2)_1$, $c \neq 0$.

If we count precisely, we see that S actually has size $(3/2)m-1$ when m is even and size $(3/2)m-1/2$ when m is odd. We will now show these upper bounds are optimal by showing that $(3/2)m-1$ is a lower bound.

So suppose $S \subset \mathbb{Z}_m^2$ blocks all noncontractible cycles. To block all $(1,0)$-cycles S must contain some (a,b)-cycle, C, in $(\mathbb{Z}_m^2)_\infty$ with $b \neq 0$. If either $|a|$ or $|b|$ is at least 2 then C contains at least $2m$ points. So we may assume that C is of type either $(0,1)$ or $(1,1)$. But now to block all cycles in $(\mathbb{Z}_m^2)_1$ that are parallel to C (i.e., have the same type as C), S must contain some other nontrivial cycle C' in $(\mathbb{Z}_m^2)_\infty$ not parallel to C. Hence we can conclude without loss of generality that one of the following three cases occurs in $(\mathbb{Z}_m^2)_\infty$: (i) S has a $(1,0)$-cycle and a $(0,1)$-cycle; (ii) S has a $(1,0)$-cycle and a $(1,1)$-cycle; or, (iii) S has a $(1,1)$-cycle and a $(1,-1)$-cycle.

For case (i), let C be the $(1,0)$-cycle and C' the $(0,1)$-cycle. Suppose that C contains t steps with vertical displacement of 1. Then it must also contain exactly t steps with vertical displacement -1, because its type is $(1,0)$. Thus C has length at least $\max(m,2t)$. Also, C is contained in the union of $t+1$ horizontal lines, so it follows that C' must have at least $m-t-1$ points not in C, since it has type $(0,1)$. Thus S has size at least $\max(m,2t)+m-t-1$, which is at least $(3/2)m-1$, as claimed.

The argument for case (ii) is identical. For case (iii) things are even easier. In this case let C be the $(1,1)$-cycle, and note that C travels up at least m steps and right at least m steps. If C is to have fewer than $(3/2)m$ points by itself, then at least $m/2$ of these steps must be shared; i.e., C must have at least $m/2$ $(1,1)$-steps. Now let C' be the $(1,-1)$-cycle. Then C' needs to take at least m steps that are either horizontal, vertical, or $(1,-1)$-steps. Since none of these are the $m/2$ $(1,1)$-steps of C, we conclude that C and C' together have at least $(3/2)m$ vertices, as claimed.

2.2 Proof of Theorem 1

Having analyzed the case of $d=2$, we will prove Theorem 1 by generalizing the example from the previous subsection to higher dimensions. Our proof uses the foam perspective described in Section 1. That is, we show a set that generates a discrete foam with periodicity $m \cdot \mathbb{Z}^d$ and has the size claimed in the theorem. To define the discrete foam boundary, it will help to first define a continuous foam.

We define inductively a set $B(r)$ in Euclidean space $\Re^d$, where $d = 2^r$. The set $B(0)$ will be the set of all $x_1 \in \Re^1$ satisfying

$$0 \leq x_1 < m.$$

In other words, $B(0) = [0,m)$. The inductive definition of $B(r) \subset \Re^d$ is

$$(x_1,\ldots,x_d) \in B(r) \quad \Leftrightarrow \quad \begin{array}{l} (x_1+x_2,\ldots,x_{d-1}+x_d) \in B(r-1) \quad \text{and} \\ (\frac{x_1-x_2}{2}, \quad \ldots, \frac{x_{d-1}-x_d}{2}) \quad \in B(r-1). \end{array}$$

Thus we have that $B(1) \subset \Re^2$ is the set of points (x_1, x_2) satisfying

$$0 \leq x_1 + x_2 < m$$

$$0 \leq x_1 - x_2 < 2m,$$

and $B(2) \subset \Re^4$ is the the set of points (x_1, x_2, x_3, x_4) satisfying

$$0 \leq x_1 + x_2 + x_3 + x_4 < m$$

$$0 \leq x_1 - x_2 + x_3 - x_4 < 2m$$

$$0 \leq x_1 + x_2 - x_3 - x_4 < 2m$$

$$0 \leq x_1 - x_2 - x_3 + x_4 < 4m,$$

and it can easily be checked that $B(r)$ is the set of points $x \in \Re^d$ satisfying $0 \leq H_r x < m u_r$, where H_r denotes the standard $2^r \times 2^r$ Hadamard matrix and u_r denotes the rth tensor power of the vector $(1, 2)$.

Let us also introduce the following notation: Let $L(r)$ denote the "lower boundary" of $B(r)$, containing all the points in $B(r)$ for which one of the inequalities hold as an equality; and, let $\overline{B}(r)$ be the closure of $B(r)$, which can also be obtained by replacing all strict inequalities by non-strict inequalities.

Since the Hadarmard matrix is orthogonal, it is easy to see that $\overline{B}(r)$ is a closed rectangular box in $\Re^d$ (although it is not axis-parallel). We will show that $B(r)$ tiles $\Re^d$ with periodicity $m \cdot \mathbb{Z}^d$. This is a consequence of the following two propositions:

Proposition 2.1. *No two points of $B(r)$ are the same modulo $m \cdot \mathbb{Z}^d$.*

Proposition 2.2. *The volume of $B(r)$ is m^d.*

Proof. (Proposition 2.1.) The proof is by induction; the statement is clearly true for $r = 0$. For larger r, suppose x is in $B(r)$ and $x + m \cdot (a_1, \ldots, a_d)$ is also in $B(r)$, where the a_i's are integers. We wish to show that all a_i's equal 0. By definition, we know that

$$(x_1 + x_2, \ldots, x_{d-1} + x_d) \in B(r-1),$$

$$(x_1 + x_2 + m \cdot (a_1 + a_2), \ldots, x_{d-1} + x_d + m \cdot (a_{d-1} + a_d)) \in B(r-1).$$

By induction, then, we get

$$a_1 + a_2 = \cdots = a_{d-1} + a_d = 0. \tag{1}$$

It follows that $a_1 - a_2, \ldots, a_{d-1} - a_d$ are all even and thus $(a_1 - a_2)/2, \ldots, (a_{d-1} - a_d)/2$ are all integers. But by definition we also know that

$$(\tfrac{x_1 - x_2}{2}, \ldots, \tfrac{x_{d-1} - x_d}{2}) \in B(r-1),$$

$$(\tfrac{x_1 - x_2}{2} + m\tfrac{a_1 - a_2}{2}, \ldots, \tfrac{x_{d-1} - x_d}{2} + m\tfrac{a_{d-1} - a_d}{2}) \in B(r-1),$$

so by induction,

$$(a_1 - a_2)/2 = \cdots = (a_{d-1} - a_d)/2 = 0. \tag{2}$$

Combining (1) and (2) we get that all a_i's are 0. This completes the induction. □

Proof. (Proposition 2.2.) As mentioned, $B(r)$ is a rectangular box, so its volume is simply the product of its side lengths. The normal vectors to its sides are the rows of the Hadamard matrix H_r, which have length $\sqrt{d}$. Thus $B(r)$'s sides have length $(m/\sqrt{d}) \cdot (u_r)_1, \ldots, (m/\sqrt{d}) \cdot (u_r)_d$, where we recall the vector u_r is the rth tensor product of $(1, 2)$. So to complete the proof it suffices to show that $\prod_{i=1}^{d} (u_r)_i = d^{d/2}$. This follows by induction since it is easy to see we have the recurrence $u_0 = 1$, $\prod_{i=1}^{d} (u_r)_i = 2^{d/2} (\prod_{i=1}^{d/2} (u_{r-1})_i)^2$. □

We have now shown that $B(r)$ tiles $\Re^d$ with periodicity $m \cdot \mathbb{Z}^d$. It follows easily that $L(r)$ generates a continious closed foam in $\Re^d$ with preriodicity $m \cdot \mathbb{Z}^d$.

Let us now return to the discrete problem in which we are interested. A natural approach would be to show that $L(r) \cap \mathbb{Z}^d$ generates a discrete foam in $(\mathbb{Z}^d)_1$ with periodicity $m \cdot \mathbb{Z}^d$, which it indeed does, and to upper-bound $S_1(m, d)$ by counting the lattice points on $L(r)$. However, to avoid the need to approximate the number of lattice points on $L(r)$, we take a slightly different tack.

Let $L'(r)$ denote a thickening of $L(r)$ to width $1/\sqrt{d}$; in other words, $L'(r) = \{x \in B(r) : \text{dist}(x, L(r)) \leq 1/\sqrt{d}\}$. Note that $L(r) + v$ generates a continuous foam in $\Re^d$ with periodicity $m \cdot \mathbb{Z}^d$ for any vector $v \in \Re^d$. From this it's easy to see that $(L'(r) + v) \cap \mathbb{Z}^d$ generates a discrete foam in $(\mathbb{Z}^d)_1$ with periodicity $m \cdot \mathbb{Z}^d$; the reason is that the normals to the faces of $L(r)$ are of the form $(\pm 1, \pm 1, \ldots, \pm 1)$, and so every edge of $(\mathbb{Z}^d)_1$ travels length at most $1/\sqrt{d}$ perpendicular to $L(r)$'s faces. Thus any infinite path in $(\mathbb{Z}^d)_1$ would have to pass through $L'(r)$.

We can now upper-bound $S_1(m, d)$ by counting the number of points in $(L'(r) + v) \cap \mathbb{Z}^d$ for any particular v. By volume considerations, it is clear that there exists a vector v such that

$$\#((L'(r) + v) \cap \mathbb{Z}^d) \leq \text{vol}(L'(r)) \leq \text{area}(L(r))/\sqrt{d}.$$

Thus to prove Theorem 1 it suffices to show that the surface area of $L(r)$ is at most $d^{\log_2(3/2)} m^{d-1} \cdot \sqrt{d}$. Since $B(r)$ is a rectangular box, the surface area of $L(r)$ is equal to the sum of the reciprocals of $B(r)$'s side lengths times its volume (i.e., m^d, by Proposition 2.2). $B(r)$'s side lengths equal $(m/\sqrt{d}) \cdot (u_r)_i$, as was mentioned in the proof of Proposition 2.2, where the vector u_r is the rth tensor power of $(1, 2)$. Thus to complete the proof we need to show that $\sum_{i=1}^{d} 1/(u_r)_i = d^{\log_2(3/2)} = (3/2)^r$. This can be proven by induction, as one can easily derive the recurrence $1/u_0 = 1$, $\sum_{i=1}^{d} 1/(u_r)_i = (3/2) \sum_{i=1}^{d/2} 1/(u_{r-1})_i$.

3 The lower bound on $S_\infty(m, d)$

In this section we prove Theorem 2. Our proof begins with the same strategy used by Saks et al. in [8], which involves *sections* and *tubes*.

Definition 4 (sections and tubes). *Given a direction $i \in [d]$ and a point $x \in \mathbb{Z}_m^d$, we define the* section based at x and perpendicular to direction i *to be the $(d-1)$-dimensional hypercube containing the points*

$$\{x + f : f \in \{0, 1\}^{i-1} \times \{0\} \times \{0, 1\}^{d-i}\}.$$

A tube in direction i *is the union of a section perpendicular to direction i with all of its translates by multiples of the vector $e_i = (0, \ldots, 0, 1, 0, \ldots, 0)$. A tube is therefore a union of m parallel sections.*

The lower bound of Saks et al., as well as our tight lower bound, is based on the following observation:

Observation 3.1. *If S is any set of vertices in $(\mathbb{Z}_m^d)_\infty$ that touches all noncontractible cycles, then S must contain at least one complete section from* every *tube.*

The proof of this observation is clear: if there were some tube for which every section had a vertex missed by S, then these vertices would form a noncontractible cycle, since all pairs of consecutive sections are completely mutually connected in $(\mathbb{Z}_m^d)_\infty$.

Given the observation above, we will now prove a lower bound of $m^d - (m-1)^d$ on the size of any subset S that contains a full section in every tube. In fact it suffices to forget about the tubes which "wrap around" the torus and think instead of the graph $([m]^d)_\infty$, which only contains the $d(m-1)^{d-1}$ tubes that are inside the grid. We prove the lower bound for any $S \subseteq [m]^d$ which contains a complete section from each one of these tubes.

The proof of Saks et al. showed that any $S \subseteq [m]^d$ containing at least one full section in each of these tubes contains at least $d(m-1)^{d-1}$ points. Their proof used a linear algebraic argument; it considered the dimension of the space spanned by indicators of the sections contained in S. We provide a more elementary argument, which gives a tight lower bound and seems to have more potential for generalizations. In particular, we would like to generalize the lower bound to the case where S is only known to contain a fixed fraction of the points of one section per tube. A good lower bound in this regime would translate to an advancement in the parallel repetition problem discussed briefly in Section 1.

Our proof goes by induction, where the key is to take a stronger induction statement. For this purpose, we define a *cube* to be a set of the form

$$\{x + f : f \in \{0,1\}^d\} \subseteq [m]^d;$$

in other words, a cube is the union of two consecutive sections. Theorem 2 follows immediately from the following:

Theorem 3. *Let S be a subset of the vertices of $[m]^d$ containing at least one complete section per tube and also containing at least c cubes. Then the cardinality of S is at least $m^d - (m-1)^d + c$.*

Proof. Let us first argue about the case $d = 2$ and $c = 0$. In this case we are considering the two-dimensional grid $[m]^2$. Tubes can be thought of as the $m-1$ vertical columns between the vertices and the $m-1$ horizontal rows between the vertices; sections can be thought of as horizontal edges and vertical edges (more accurately, as the pair of vertices making up these edges). Suppose S contains at least one horizontal edge per column and one vertical edge per row. When taken together, its clear that these $2m-2$ edges cannot form any cycle since they never have two edges "one above the other" (or "one to the left of the other"). Since an acyclic graph with $2m-2$ has exactly $2m-1 = m^2 - (m-1)^2$ vertices, the proof of the $d = 2$, $c = 0$ case is complete.

We next consider the $d = 2$ case for general c. In this case, we know that S contains at least $m-1$ vertical edges (sections) and it is clear that it must contain at least c more vertical edges because of the presence of c cubes (cubes are squares, in two dimensions). We have so far identified

$m-1+c$ vertical edges contained in S. Now consider adding the $m-1$ horizontal sections that S must contain. The resulting set of $2m-2+c$ edges must still be acyclic since it has no two horizontal edges in the same tube. Thus it contains $2m-1+c = m^2-(m-1)^2+c$ vertices as required by the induction.

With the case $d=2$ completely proven, we move to the induction on the dimension d. So suppose S is a subset of $[m]^d$ with at least one section per tube and also at least c cubes. Consider the set of sections perpendicular to the dth direction. We know that there are at least $(m-1)^{d-1}$ of these which are contained in S — one per tube going the dth direction. There must also be at least c tubes in the dth direction where S contains an additional section, because of the c cubes it contains. Let us stratify these sections according to what level $1,\ldots,m$ they are on in the dth direction. Specifically, say we have c_i of them on level i, where $c_1+\cdots+c_m \geq (m-1)^{d-1}+c$.

We now view the ith level as an inductive instance in dimension $d-1$. Because S has at least one section per tube in $[m]^d$, it is easy to see that it also has at least one (lower-dimensional) section per (lower-dimensional) tube in $[m]^{d-1}$. It also has at least c_i cubes. So by induction, S has at least $m^{d-1}-(m-1)^{d-1}+c_i$ vertices on the ith level of $[m]^d$. Summing this over i yields at least

$$m(m^{d-1}-(m-1)^{d-1})+(m-1)^{d-1}+c = m^d-(m-1)^d+c$$

as a lower bound for the number of points in S. □

Theorem 2 follows from Theorem 3 by taking $c=0$.

References

[1] M. C. Azizoğlu and Ö. Eğecioğlu. The isoperimetric number of d-dimensional k-ary arrays. *International Journal of Foundations of Computer Science*, 10(3):289–300, 1999.

[2] B. Bollobás and I. Leader. An isoperimetric inequality on the discrete torus. *SIAM J. Discret. Math.*, 3(1):32–37, 1990.

[3] F. R. K. Chung and P. Tetali. Isoperimetric inequalities for cartesian products of graphs. *Combinatorics, Probability & Computing*, 7(2):141–148, 1998.

[4] U. Feige. Error reduction by parallel repetition — the state of the art. Technical report, Weizmann Institute of Science, 1995.

[5] U. Feige. Personal communication. 2005.

[6] R. Raz. A parallel repetition theorem. *SIAM Journal on Computing*, 27(3):763–803, June 1998.

[7] A. Rós. The isoperimetric problem. Lecture notes. MSRI summer school, 2001. `http://www.ugr.es/~aros/isoper.htm`.

[8] M. E. Saks, A. Samorodnitsky, and L. Zosin. A lower bound on the integrality gap for minimum multicut in directed networks. *Combinatorica*, 24(3):525–530, 2004.