



Understanding Parallel Repetition Requires Understanding Foams

Uriel Feige

Weizmann Institute & Microsoft Research
uriel.feige@microsoft.com

Guy Kindler*

Weizmann Institute
guy.kindler@weizmann.ac.il

Ryan O'Donnell*

Carnegie Mellon University
odonnell@cs.cmu.edu

May 7, 2007

Abstract

Motivated by the study of Parallel Repetition and also by the Unique Games Conjecture, we investigate the value of the “Odd Cycle Games” under parallel repetition. Using tools from discrete harmonic analysis, we show that after d rounds on the cycle of length m , the value of the game is at most $1 - (1/m) \cdot \tilde{\Omega}(\sqrt{d})$ (for $d \leq m^2$, say). This beats the natural barrier of $1 - \Theta(1/m)^2 \cdot d$ for Raz-style proofs [Raz98, Hol06] (see [Fei95]) and also the SDP bound of Feige-Lovász [FL92, GW95]; however, it just barely fails to have implications for Unique Games. On the other hand, we also show that improving our bound would require proving nontrivial lower bounds on the surface area of high-dimensional foams. Specifically, one would need to answer: What is the least surface area of a cell that tiles \mathbb{R}^d by the lattice \mathbb{Z}^d ?

1 Introduction

In this paper we observe some intriguing relations between problems from different areas, which seem unrelated at first look. In particular, we show that any significant improvement on bounds for parallel repetition on 2-player games would entail progress on a certain foam problem in \mathbb{R}^d which is currently wide open. On the other hand, we manage to translate a known bound for the foam problem to get an improved bound for a special case of the Strong Parallel Repetition Problem.

1.1 Problems and relations

We now state some of the problems which are discussed in this paper, and discuss their relations in detail. The definitions of other problems and of some of the notions used here appear in Subsection 1.2 below.

We begin with the Strong Parallel Repetition Problem, which asks whether the parameters in the Parallel Repetition Theorem [Raz98, Hol06] can be significantly improved. It is a natural open question in game theory, and as discussed in Section 2, it has implications in complexity theory as well. This problem was one of the original motivations of this work.

Strong Parallel Repetition Problem: Let \mathfrak{G} be a 2-player 1-round game where the players' answers are in the sets \mathcal{A} and \mathcal{B} respectively. Is it true that $\text{val}(\mathfrak{G}) \leq 1 - \epsilon$ implies $\text{val}(\mathfrak{G}^{\otimes d}) \leq (1 - \Omega(\epsilon))^{d/\log(|\mathcal{A}||\mathcal{B}|)}$?

The Strong Parallel Repetition Problem for XOR Games (defined in Subsection 1.2) is an interesting special case of the Strong Parallel Repetition Problem, and a very special case of that is our Odd Cycle Parallel Repetition Problem, which asks whether one can get improved parameters for parallel repetition over a very specialized type of games.

*Some of this research was done while the author was at Microsoft Research. Research was partially funded by the Israel Science Foundation (ISF).

Odd Cycle Parallel Repetition Problem: Let \mathfrak{G}_{C_m} be the length- m Odd Cycle Game (see Definition 1.3), satisfying $\text{val}(\mathfrak{G}_{C_m}) = 1 - \Theta(1/m)$. What is $\text{val}(\mathfrak{G}_{C_m}^{\otimes d})$, its value after d rounds of parallel repetition? Is it true that $\text{val}(\mathfrak{G}_{C_m}^{\otimes d}) \leq 1 - (1/m) \cdot \Omega(d)$ for $d \leq m$?

The observation which is reflected in the title of this paper is that upper bounds for the value of $\mathfrak{G}_{C_m}^{\otimes d}$ translate to lower bounds on the surface area of cells that tile \mathbb{R}^d by a certain lattice. Thus good bounds for the Strong Parallel Repetition Problem, which must apply to the Odd Cycle Parallel Repetition subproblem, also translate to bounds for the surface area of certain periodic foams, and thus must increase our understanding of these little understood objects – many problems of minimal surface-area foams are longstanding, notorious problems in geometry, and may be difficult to solve.

Foam Problem on $\mathbb{R}^d/\mathbb{Z}^d$: What is the least surface area $A(d)$ of a cell that tiles \mathbb{R}^d by \mathbb{Z}^d ?

Bounds for the Odd Cycle Parallel Repetition Problem translate to foams through a problem of eliminating cycles in a certain torus-like graph.

Cycle Elimination Problem on $(\mathbb{Z}_m^d)_\infty$: Let $(\mathbb{Z}_m^d)_\infty$ denote the discrete torus graph on vertex set \mathbb{Z}_m^d in which two vertices are connected if their ℓ_∞ distance is at most 1. What is the least fraction of edges $\delta(d, m)$ that need to be deleted from $(\mathbb{Z}_m^d)_\infty$ so as to eliminate all of the topologically nontrivial cycles?

The Cycle Elimination Problem on $(\mathbb{Z}_m^d)_\infty$ is in fact almost equivalent to the Odd Cycle Parallel Repetition Problem (equivalence truly holds for a slight variation of it, the Odd Cycle Elimination Problem on $K_2 \times (\mathbb{Z}_m^d)_\infty$). In the most technically challenging part of this work, we manage to translate a simple lower-bound for the Foam Problem back to a bound for the Cycle Elimination Problem, and from it, to the Odd Cycle Parallel Repetition Problem, getting bounds that are better than can be achieved with the current Parallel Repetition Theorems.

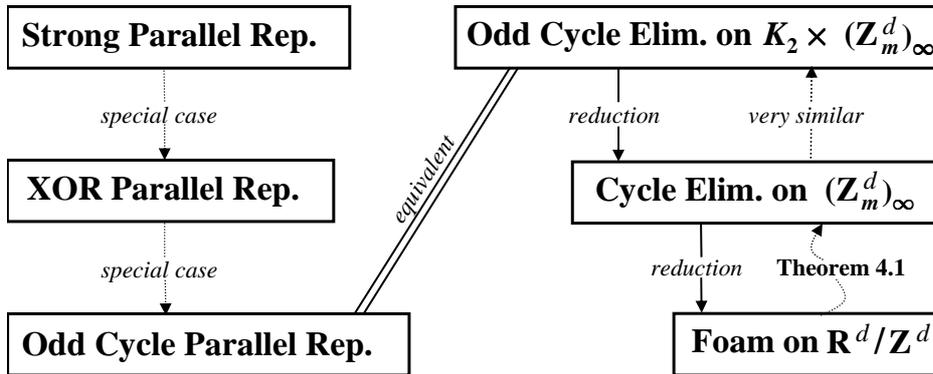


Figure 1: The *Strong Parallel Repetition Problem for XOR Games* is special case of the *Strong Parallel Repetition Problem*. A very special case of that is our *Odd Cycle Parallel Repetition Problem*. This problem is equivalent to the *Odd Cycle Elimination Problem on $K_2 \times (\mathbb{Z}_m^d)_\infty$* , a slight variant of our *Cycle Elimination Problem on $(\mathbb{Z}_m^d)_\infty$* . When $d \ll m$ this last problem seems essentially similar to the *Foam Problem on $\mathbb{R}^d/\mathbb{Z}^d$* , which is connected to longstanding, notorious problems in geometry.

1.2 Definitions

Before we give precise statements of our results, let us give the definitions of the notions used so far, and of a few more problems. We begin with 2-player 1-round games, and parallel repetition.

Definition 1.1 A 2-player 1-round (2P1R) game \mathfrak{G} with question sets \mathcal{X} and \mathcal{Y} and answer sets \mathcal{A} and \mathcal{B} is a tuple that consists of a probability distribution π on $\mathcal{X} \times \mathcal{Y}$ and a family of predicates $(V_{xy})_{x \in \mathcal{X}, y \in \mathcal{Y}}$ each on $\mathcal{A} \times \mathcal{B}$. It describes a game in which a “verifier” chooses (X, Y) randomly according to π and sends

question X to player “Alice” and question Y to player “Bob”. Without further communication, Alice must reply with an answer A and Bob with an answer B . The verifier then checks whether $V_{XY}(A, B)$ holds; in this case Alice and Bob “win”.

The value of the game is

$$\text{val}(\mathfrak{G}) = \sup_{f_A, f_B} \mathbf{E}_{(X, Y) \leftarrow \pi} [V_{XY}(f_A(X), f_B(Y))],$$

where the sup is over strategies $f_A : \mathcal{X} \rightarrow \mathcal{A}$ and $f_B : \mathcal{Y} \rightarrow \mathcal{B}$ for Alice and Bob.

Equivalently, one can think of a 2P1R game as a constraint satisfaction problem on a weighted bipartite graph $(\mathcal{X} \times \mathcal{Y}, \pi)$, with label sets \mathcal{A} for \mathcal{X} and \mathcal{B} for \mathcal{Y} and constraints given by the V_{xy} ’s. In this paper we will be particularly interested in some special classes of 2P1R games:

Definition 1.2 A 2P1R game is binary if $\mathcal{A} = \mathcal{B} = \{0, 1\}$. It is said to be a unique game if: (i) $|\mathcal{A}| = |\mathcal{B}|$ and (ii) for every $x \in \mathcal{X}$, $y \in \mathcal{Y}$, there is a permutation $\sigma_{xy} : \mathcal{A} \rightarrow \mathcal{B}$ such that $V_{xy}(a, b)$ is true iff $b = \sigma_{xy}(a)$. Finally, a binary unique game is called an XOR game; for such games, each constraint V_{xy} is of the form $a + b = 0 \pmod{2}$ or $a + b = 1 \pmod{2}$; equivalently, $a = b$ or $a \neq b$.

Perhaps the simplest family of XOR games is the family of Odd Cycle Games:

Definition 1.3 The m -Cycle Game (for m odd), denoted \mathfrak{G}_{C_m} , is a 2P1R game in which $\mathcal{X} = \mathcal{Y} = \mathbb{Z}_m$ and $\mathcal{A} = \mathcal{B} = \mathbb{Z}_2$. The distribution π on $\mathbb{Z}_m \times \mathbb{Z}_m$ is given by choosing the first coordinate X uniformly at random and setting the second coordinate Y to be $X + \Lambda$, where Λ is chosen independently and uniformly from $\{-1, 0, 1\}$. The winning predicate $V_{xy}(a, b)$ is defined to be the equality constraint $a = b$ when $x = y$, and to be the inequality constraint $a \neq b$ when $x \neq y$.

Remark 1.4 The m -Cycle Games are XOR games. They satisfy $\text{val}(\mathfrak{G}_{C_m}) = 1 - (2/3)(1/m)$.

Given any 2P1R game, we can produce new, more complicated 2P1R games via the operation known as parallel repetition:

Definition 1.5 Given a 2P1R game \mathfrak{G} as in Definition 1.1 and $d \in \mathbb{N}$, the d -“round” parallel repeated version is the 2P1R game $\mathfrak{G}^{\otimes d}$ with question sets \mathcal{X}^d and \mathcal{Y}^d , answer sets \mathcal{A}^d and \mathcal{B}^d , the product distribution $\pi^{\otimes d}$ on $\mathcal{X}^d \times \mathcal{Y}^d$, and the family of predicates $(V_{xy}^{\otimes d})_{x \in \mathcal{X}^d, y \in \mathcal{Y}^d}$ defined by

$$V_{xy}^{\otimes d}(a, b) = V_{x_1 y_1}(a_1, b_1) \wedge V_{x_2 y_2}(a_2, b_2) \wedge \cdots \wedge V_{x_d y_d}(a_d, b_d).$$

One of our main problems, the Odd Cycle Parallel Repetition Problem, is to give bounds on the value of \mathfrak{G}_{C_m} after d rounds of parallel repetition. When $\mathfrak{G}_{C_m}^{\otimes d}$ is viewed as a constraint satisfaction problem on a bipartite graph, the graph in question is the double cover of a discrete torus graph.

Definition 1.6 Given $d, m \in \mathbb{N}$, the discrete torus graph with ℓ_∞ edge structure, denoted $(\mathbb{Z}_m^d)_\infty$, is the graph on vertex set \mathbb{Z}_m^d in which two vertices are connected if their ℓ_∞ distance is at most 1. For the purposes of edge counting, the self-loops of $(\mathbb{Z}_m^d)_\infty$ count as half-edges.

Definition 1.7 The double cover of the discrete torus graph $(\mathbb{Z}_m^d)_\infty$, which we denote by $K_2 \times (\mathbb{Z}_m^d)_\infty$, is the bipartite graph whose vertex set consists of two disjoint copies of \mathbb{Z}_m^d , called the A-side and the B-side. The graph has an edge from x on the A-side to y on the B-side whenever (x, y) is an edge in $(\mathbb{Z}_m^d)_\infty$. Given a vertex, edge, or cycle in $K_2 \times (\mathbb{Z}_m^d)_\infty$, we say its projection to $(\mathbb{Z}_m^d)_\infty$ is the vertex, edge, or cycle in $(\mathbb{Z}_m^d)_\infty$ one gets by ignoring the distinction between the A- and B-sides.

Determining $\text{val}(\mathfrak{G}_{C_m}^{\otimes d})$ can be viewed as the problem of finding an optimal \mathbb{Z}_2^m -labeling of the vertices of $K_2 \times (\mathbb{Z}_m^d)_\infty$ so as to satisfy certain constraints. However it turns out that it is equivalent to a simpler problem. To describe this problem though, we need to introduce some topological notions.

Definition 1.8 The d -dimensional unit cubic torus, denoted \mathcal{T}^d , is $\mathbb{R}^d / \mathbb{Z}^d$ (interpreted as a Riemannian manifold, say); i.e., the unit cube $[0, 1]^d$ with opposite faces identified.

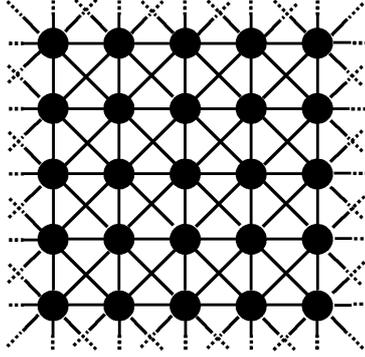


Figure 2: $(\mathbb{Z}_5^2)_\infty$ — self-loops not pictured

The graph $(\mathbb{Z}_m^d)_\infty$ can be naturally embedded in \mathcal{T}^d , with its edges being straight line segments. We sometimes think of it as a discretization of \mathcal{T}^d .

Definition 1.9 Given an (oriented) cycle L in the graph $(\mathbb{Z}_m^d)_\infty$, its homotopy class, denoted $[L] \in \mathbb{Z}^d$, is defined to be the homotopy class of the closed loop in \mathcal{T}^d with which it is naturally identified. I.e., if L “wraps around the boundary” c_i times in the i th direction, then $[L] = (c_1, c_2, \dots, c_d)$. For a cycle L in $K_2 \times (\mathbb{Z}_m^d)_\infty$ we define its homotopy class $[L]$ to be the homotopy class of its projection to $(\mathbb{Z}_m^d)_\infty$.

Definition 1.10 We say that a cycle L (in $(\mathbb{Z}_m^d)_\infty$ or in $K_2 \times (\mathbb{Z}_m^d)_\infty$) is topologically nontrivial if $[L] \neq 0$. We say that it is topologically odd if $[L]_i$ is odd for some coordinate $1 \leq i \leq d$. Note that a topologically odd cycle is topologically nontrivial.

We now present the simpler problem which is equivalent to determining $\text{val}(\mathfrak{G}_{C_m}^{\otimes d})$:

Odd Cycle Elimination Problem on $K_2 \times (\mathbb{Z}_m^d)_\infty$: What is the least fraction of edges $\delta'(d, m)$ that need to be deleted from $K_2 \times (\mathbb{Z}_m^d)_\infty$ so as to eliminate all of the topologically odd cycles?

Finally, we mention the notion of foams. In geometric measure theory, a *foam* is a “a piecewise smooth arrangement of connected, bounded regions (‘bubbles’) tiling space (of arbitrary dimension) [in which] the bubble volumes are taken to be fixed.” [WKC+04]. Usually one is interested in studying foams of minimal surface area. In particular, the Foam Problem on $\mathbb{R}^d/\mathbb{Z}^d$ is concerned with finding the best foam among those that are periodic with respect to the integer lattice.

1.3 Our results

We now describe our results; specifically, our bounds for the Odd Cycle Parallel Repetition Problem. Referring back to Figure 1, the reductions illustrated there imply that an upper bound for the Foam Problem on $\mathbb{R}^d/\mathbb{Z}^d$ — i.e., a cell with small surface area that tiles \mathbb{R}^d by \mathbb{Z}^d — yields good strategies for the players in $\mathfrak{G}_{C_m}^{\otimes d}$. Specifically, we have the following results, the proofs of which are not hard:

Theorem 1.11

$$\begin{aligned} \text{val}(\mathfrak{G}_{C_m}^{\otimes d}) &= 1 - \delta'(d, m) \\ &\quad (\text{proven in Proposition 2.4}) \\ &\geq 1 - 2\delta(d, m) \\ &\quad (\text{proven in Remark 2.5}) \\ &\geq 1 - (1/m) \cdot O(A(d)) \\ &\quad (\text{proven in Theorem 3.1}). \end{aligned}$$

Regarding $A(d)$, the least surface area of a cell that tiles d -dimensional space by the integer lattice, in Section 3.2 we remark that the following easy bounds hold (it turns out that $A(d)/2$ is a slightly more natural quantity to study than $A(d)$):

Proposition 1.12

$$d - \exp(-O(d \log d)) \geq \frac{A(d)}{2} \geq \begin{cases} \sqrt{\pi e/2} \sqrt{d} - o(\sqrt{d}), \\ \sqrt{d}. \end{cases}$$

The first lower bound is in fact always better than the second, although just by the constant factor $\sqrt{\pi e/2} \approx 2.066$. However it uses the classical isoperimetric inequality, whereas the \sqrt{d} bound just uses elementary reasoning.

If we place a cell that tiles \mathbb{R}^d by \mathbb{Z}^d inside the unit torus \mathcal{T}^d , its surface has the property that it blocks all homotopically nontrivial loops in \mathcal{T}^d . Thus $A(d)/2$ can also be viewed as the least amount of surface that needs to be deleted from \mathcal{T}^d so as to eliminate all topologically nontrivial cycles from \mathcal{T}^d (one divides by two since each point on the cell's surface is counted twice — from the inside and from the outside). This forms the intuitive connection between our Foam Problem and our Cycle Elimination Problem, especially for the case of large m in which the fine granularity of $(\mathbb{Z}_m^d)_\infty$ approximates the continuous torus \mathcal{T}^d .

Unfortunately, it's not clear that the connection between the two problems can be put precisely; in contrast to Theorem 3.1 we don't show that lower bounds for $A(d)$ automatically translate into lower bounds for $\delta(d, m)$. The main issue is that whereas solutions to the Foam Problem are *defined* to be surfaces, it's not clear how to show that the minimal solution to the Cycle Elimination Problem on $(\mathbb{Z}_m^d)_\infty$ is “surface-like”. In fact, it seems that it won't be if m is small compared to d .

Nevertheless, the main theorem of this paper shows that the elementary reasoning which shows $A(d)/2 \geq \sqrt{d}$ can be translated — albeit very elaborately — into an essentially equivalent lower bound on $\delta(d, m)$:

Main Theorem, Theorem 4.1 *Provided $m \geq \sqrt{d/\log d}$,*

$$\delta(d, m) \geq (1/m) \cdot \Omega(\sqrt{d/\log d}),$$

and indeed $\delta'(d, m) \geq (1/m) \cdot \Omega(\sqrt{d/\log d})$ as well.

Using $\text{val}(\mathfrak{G}_{C_m}^{\otimes d}) = 1 - \delta'(d, m)$ (proven in Proposition 2.4 as mentioned in Theorem 1.11) we get as a corollary:

Corollary 1.13 *Provided $d \leq m^2 \log m$,*

$$\text{val}(\mathfrak{G}_{C_m}^{\otimes d}) \leq 1 - (1/m) \cdot \Omega(\sqrt{d/\log d}).$$

Unfortunately, this result does not get as far as we wanted for the Odd Cycle Parallel Repetition Problem. However from Theorem 1.11 we see that getting an improved lower bound (beyond the $\sqrt{\log d}$ factor) would require making an advancement on the Foam Problem on $\mathbb{R}^d/\mathbb{Z}^d$, specifically, proving a better lower bound than $\Theta(\sqrt{d})$.

Finally, we note that *perfect parallel repetition* does not hold for m -cycle games. That is, while the value of the m -cycle game is $\text{val}(\mathfrak{G}_{C_m}) = 1 - (2/3)(1/m)$ (see Remark 1.4), there are at least some values of d for which $\text{val}(\mathfrak{G}_{C_m}^{\otimes d}) > (1 - (2/3)(1/m))^d$.

Theorem 1.14 *There exists a constant $0 < \delta < 1$, such that for sufficiently large odd m and for the largest integer d such that $(1 - 2/3m)^d > 1/m$,*

$$\text{val}(\mathfrak{G}_{C_m}^{\otimes d}) > 1/m^\delta.$$

Due to space limitations, the proof of Theorem 1.14 will only appear in the full version of the paper.

1.4 Outline of the rest of paper

In Section 2 we discuss Parallel Repetition in general and explain some motivations for studying it on the Odd Cycle Games — namely, understanding the rate of parallel repetition, and the Unique Games Conjecture. We also summarize what is known about the Odd Cycle Parallel Repetition Problem and prove Proposition 2.4. In Section 3 we discuss foam problems and prove Theorem 3.1, and in Section 4 we prove our main theorem, Theorem 4.1.

1.5 Acknowledgments

The authors are grateful to John M. Sullivan for his expertise on foams, for directing us to [Cho89], and for his suggestion to calculate the surface area of the Voronoi cell of the A_d^* lattice. The authors would also like to thank Madhu Sudan and Avi Wigderson for helpful discussions.

2 Parallel Repetition

Two-prover one-round games were introduced to complexity theory in the study of interactive proofs [BOGKW88]. In this context, $\text{val}(\mathfrak{G})$ often represents the probability that the proof of a false statement is accepted. Of course if $\text{val}(\mathfrak{G})$ is close to 1, it is important to analyze how \mathfrak{G} can be transformed into a new 2P1R game with much smaller value. Parallel repetition is the most natural thing to try.

By using the optimal strategies for \mathfrak{G} in a coordinatewise fashion, it is easy to see that for any game \mathfrak{G} it holds that $\text{val}(\mathfrak{G}^{\otimes d}) \geq \text{val}(\mathfrak{G})^d$. But early on it was observed [For89] that the reverse inequality need not be true, since optimal strategies $f_A : \mathcal{X}^d \rightarrow \mathcal{A}^d$ and $f_B : \mathcal{Y}^d \rightarrow \mathcal{B}^d$ for $\mathfrak{G}^{\otimes d}$ are not necessarily coordinatewise. It is not particularly easy even to improve on the trivial bound $\text{val}(\mathfrak{G}^{\otimes d}) \leq \text{val}(\mathfrak{G})$ (which holds because if the players are to win in the repeated game they must at least win on the first coordinate). In fact, no improvement on this is even possible if d is only 2, as Feige [Fei91] showed for a certain 2P1R game with answer size 3.

Extra impetus for understanding the rate of parallel repetition came from the fundamental PCP Theorem [AS98, ALM⁺98] of computational complexity, which is equivalent to the following statement: There are universal constants $k \in \mathbb{N}$ and $\epsilon > 0$ such that given a 2P1R game \mathfrak{G} with answer sets of size k , it is NP-hard to distinguish the cases $\text{val}(\mathfrak{G}) = 1$ and $\text{val}(\mathfrak{G}) \leq 1 - \epsilon$. For applications to hardness of finding approximate solutions to combinatorial optimization problems, it is of tremendous importance to improve this to hardness of distinguishing $\text{val}(\mathfrak{G}) = 1$ from $\text{val}(\mathfrak{G}) \leq \epsilon$ for every $\epsilon > 0$, with k being allowed to depend on ϵ . This was first achieved by Feige and Kilian [FK00]. However a bound with much better dependence of k on ϵ followed immediately when Raz proved his famous Parallel Repetition Theorem:

Theorem 2.1 (*Raz’s Parallel Repetition Theorem.*) *For every $v < 1$ there exists $v' < 1$ such that the following holds: If \mathfrak{G} is any 2P1R game with $\text{val}(\mathfrak{G}) < v$, then*

$$\text{val}(\mathfrak{G}^{\otimes d}) \leq (v')^{d/\log(|\mathcal{A}||\mathcal{B}|)}.$$

As was necessary for PCP applications, there is no dependence here on the sizes of the questions sets, $|\mathcal{X}|$ and $|\mathcal{Y}|$. Regarding the sizes of the answer sets, Feige and Verbitsky [FV02] demonstrated that the logarithmic dependence is essentially necessary. However another quantitative aspect of the theorem has remained relatively unexplored: the dependence of v' on v in the range near 1.

2.1 The rate of parallel repetition

In Raz’s original work, when v is of the form $1 - \epsilon$ for small ϵ , the quantity v' is equal to $1 - \text{poly}(\epsilon)$. Holenstein [Hol06] recently tightened this dependence to $1 - \Omega(\epsilon^3)$:

Theorem 2.2 (*Holenstein.*)

$$\text{val}(\mathfrak{G}) \leq 1 - \epsilon \quad \Rightarrow \quad \text{val}(\mathfrak{G}^{\otimes d}) \leq (1 - \Omega(\epsilon^3))^{d/\log(|\mathcal{A}||\mathcal{B}|)}.$$

For some 2P1R games, an early result of Feige and Lovász [FL92] may give an improved bound. In particular, for each game \mathfrak{G} , Feige and Lovász introduced a certain semidefinite program whose value, $\bar{\sigma}(\mathfrak{G})$, is an upper bound on the game's value $\text{val}(\mathfrak{G})$. They then showed that $\text{val}(\mathfrak{G}^{\otimes d}) \leq \bar{\sigma}(\mathfrak{G})^d$. This bound is not always useful, since in principle it may be that $\text{val}(\mathfrak{G}) \leq 1 - \epsilon$ and yet $\bar{\sigma}(\mathfrak{G}) = 1$. Nevertheless, as describe below, there are some classes of games with $\text{val}(\mathfrak{G}) \leq 1 - \epsilon$ for which $\bar{\sigma}(\mathfrak{G})$ has a provably better dependence on ϵ than Holenstein's $1 - \Omega(\epsilon^3)$.

A natural question to ask is what the best possible dependence on ϵ can be. In particular, can it be made linear? This is the content of our Strong Parallel Repetition Problem from Section 1.1, which we repeat here:

Strong Parallel Repetition Problem:

Does $\text{val}(\mathfrak{G}) \leq 1 - \epsilon$ imply that $\text{val}(\mathfrak{G}^{\otimes d}) \leq (1 - \Omega(\epsilon))^{d/\log(|\mathcal{A}||\mathcal{B}|)}$?

Understanding this problem was one of the main motivations of this work. Some remarks on it are in order:

1. The problem is open even in the significantly special case of XOR Games; in this case, the question is simply whether $\text{val}(\mathfrak{G}) \leq 1 - \epsilon$ implies $\text{val}(\mathfrak{G}^{\otimes d}) \leq (1 - \Omega(\epsilon))^d$. This is the Strong Parallel Repetition Problem for XOR Games mentioned in Figure 1 of Section 1.1.
2. As stated in Theorem 1.14, *perfect parallel repetition* does not even hold for XOR games (an earlier example of this appears in [BCH⁺02]). On the other hand, perfect parallel repetition *does* hold for XOR games in the setting of *quantum* 2P1R games, as shown in [CSUU06]; i.e., $\text{val}_q(\mathfrak{G}^{\otimes d}) = \text{val}_q(\mathfrak{G})^d$ if \mathfrak{G} is an XOR game.
3. Feige [Fei95] sketched an argument for why the inductive proof strategy of Raz (and Holenstein) does not seem likely to be able to give a bound better than quadratic dependence, $(1 - \epsilon^2)^{d/\log(|\mathcal{A}||\mathcal{B}|)}$.
4. For XOR games \mathfrak{G} it is known (see [FG95]) that Feige and Lovász's $\bar{\sigma}(\mathfrak{G})$ is equal to Goemans and Williamson's [GW95] SDP relaxation for Max-2Lin,

$$\bar{\sigma}(\mathfrak{G}) = \sup_{\substack{u: \mathcal{X} \rightarrow S^{n-1} \\ v: \mathcal{Y} \rightarrow S^{n-1}}} \mathbf{E}_{\pi} \left[\frac{1}{2} \pm \frac{1}{2} \langle u(X), v(Y) \rangle \right], \tag{1}$$

where S^{n-1} denotes the unit sphere in n dimensions (the sup is over n as well) and the \pm sign is chosen to be $+$ when V_{XY} is the equality constraint and $-$ when it is the inequality constraint. As a consequence of Goemans and Williamson's rounding algorithm we get the following quadratic bound:

Theorem 2.3 *Let \mathfrak{G} be an XOR game. Then*

$$\text{val}(\mathfrak{G}) \leq 1 - \epsilon \quad \Rightarrow \quad \text{val}(\mathfrak{G}^{\otimes d}) \leq (1 - c \cdot \epsilon^2)^d,$$

where $c = \pi^2/4 - o_\epsilon(1)$.

5. Non-strong parallel repetition guarantees have the seeming shortcoming that they are worse than trivial for superconstant values of d . For example, even with Theorem 2.3 we get no nontrivial upper bound on $\text{val}(\mathfrak{G}^{\otimes d})$ for XOR games \mathfrak{G} having value $1 - \epsilon$ until $d \geq \Omega(1/\epsilon)$. However, we cannot rule out the possibility that this may be the true behavior for some such games.

2.2 The Unique Games Conjecture

One of the original motivations for this paper has to do with the now-notorious Unique Games Conjecture of Khot [Kho02]:

Unique Games Conjecture *For every $\epsilon > 0$ there exists $k \in \mathbb{N}$ such that given a unique 2P1R game \mathfrak{G} with answer sets of size k , it is NP-hard to distinguish the cases $\text{val}(\mathfrak{G}) \geq 1 - \epsilon$ and $\text{val}(\mathfrak{G}) \leq \epsilon$.*

The Unique Games Conjecture has been shown to have strong consequences for hardness of approximation [CKK⁺05, DMR06, Kho02, KKMO04, KO06, KR03, KV05, ST06], and by now it is becoming common to refer to these as “UG-hardness” results. This raises a natural question: are there any interesting “UG-completeness” results? I.e., which of the consequences of the UGC are actually equivalent to it?

The most intriguing possibility comes from the Max-Cut problem. Khot, Kindler, Mossel and O’Donnell [KKMO04], in combination with a result of Mossel, O’Donnell and Oleszkiewicz [MOO05], proved the following UG-hardness result: for sufficiently small $\epsilon > 0$, it is UG-hard to distinguish graphs in which the maximum cut is a $1 - \epsilon$ fraction of the edges from graphs in which the maximum cut is a $1 - \sqrt{2/\pi} \cdot \sqrt{\epsilon}$ fraction of the edges. (No stronger hardness result is possible, because of the Goemans-Williamson SDP-rounding result that gives us Theorem 2.3.) It would be particularly appealing to show the reverse, that the UGC follows from $1 - \epsilon$ vs. $1 - \Omega(\sqrt{\epsilon})$ NP-hardness for Max-Cut: Max-Cut is the simplest APX-hard problem and perhaps one could even conceive of improving the best known hardness result for Max-Cut ($1 - \epsilon$ vs. $1 - (5/4)\epsilon$, due to [Hås01, TSSW00]) by gadgets.

Since Max-Cut is a constraint satisfaction problem with “unique constraints”, and since parallel repetition preserves the uniqueness property, it seems highly natural to try to use straight parallel repetition to prove the UGC given $1 - \epsilon$ vs. $1 - \Omega(\sqrt{\epsilon})$ hardness for Max-Cut. There is a slight technicality to deal with first: hard Max-Cut instances are of course not bipartite, so we would not be in the setting of 2P1R games. However it is easy to convert a Max-Cut instance into an XOR 2P1R game: with some constant probability have the verifier pick a random vertex, ask both Alice and Bob for its label, and check that the answers are the same; with the remaining probability have the verifier pick a random edge, ask Alice for one endpoint’s label and Bob for the other’s, and check that the answers are different. Using this construction it is easy to see that $1 - \epsilon$ vs. $1 - \Omega(\sqrt{\epsilon})$ hardness for Max-Cut implies hardness of distinguishing $\text{val}(\mathfrak{G}) \geq 1 - \epsilon$ from $\text{val}(\mathfrak{G}) \leq 1 - \Omega(\sqrt{\epsilon})$ for XOR games.

Rather, the real catch with this approach is that non-strong parallel repetition may destroy a $1 - \epsilon$ vs. $1 - \Omega(\sqrt{\epsilon})$ gap. In other words, suppose we try to reduce an XOR game \mathfrak{G} to the unique game $\mathfrak{G}^{\otimes d}$ with answer size 2^d , perhaps with d taken to be $1/(\epsilon \log(1/\epsilon))$. In the “completeness case” when $\text{val}(\mathfrak{G}) \geq 1 - \epsilon$, we get $\text{val}(\mathfrak{G}^{\otimes d}) \geq (1 - \epsilon)^{1/(\epsilon \log(1/\epsilon))} \geq 1 - 1/\log(1/\epsilon)$. However, though we might hope to get $\text{val}(\mathfrak{G}^{\otimes d}) \leq (1 - \Omega(\sqrt{\epsilon}))^{1/(\epsilon \log(1/\epsilon))} = \exp(-\tilde{\Omega}(1/\sqrt{\epsilon}))$ in the “soundness case”, in fact Theorem 2.3 gives us exactly nothing, since the $1 - \Omega(\sqrt{\epsilon})$ jumps up to $1 - c \cdot (\Omega(\sqrt{\epsilon}))^2 > 1 - \epsilon$.

Thus we see that this approach will work if and only if we can break the “quadratic barrier” for parallel repetition of XOR games, found in both Feige’s argument (remark #3 of Section 2.1) and Theorem 2.3.

2.3 The Odd Cycle Games

To study the Strong Parallel Repetition Problem for XOR Games, we look at the simplest possible parameterized family of XOR games whose values approach 1: the Odd Cycle Games. Recall the definition of \mathfrak{G}_{C_m} from Definition 1.3 and that $\text{val}(\mathfrak{G}_{C_m}) = 1 - (2/3)(1/m) = 1 - \Theta(1/m)$.

The best upper bound on $\text{val}(\mathfrak{G}_{C_m}^{\otimes d})$ previously known follows from the Feige-Lovász SDP method: It is well-known that the optimal solution to the Goemans-Williamson SDP for the m -cycle (m odd) maps $x \in \mathbb{Z}_m$ to $(\cos \frac{x(m-1)\pi}{m}, \sin \frac{x(m-1)\pi}{m}) \in S^1$; it is easy to check further that the optimal solution to (1) for \mathfrak{G}_{C_m} involves taking both u and v to be this map. This implies

$$\bar{\sigma}(\mathfrak{G}_{C_m}) = \frac{2}{3} + \frac{1}{3} \cos(\pi/m) = 1 - c' \cdot (1/m)^2,$$

and thus

$$\text{val}(\mathfrak{G}_{C_m}^{\otimes d}) \leq 1 - \Theta(1/m)^2 \cdot d, \quad \text{for } d \leq m^2, \quad (2)$$

where $c' = \pi^2/6 - o_m(1)$. Notice that this improves on Theorem 2.3 by only a constant factor.

The main result of this paper improves on this bound (for $d \leq m^2/\log m$):

Corollary 1.13 *Provided $d \leq m^2 \log m$,*

$$\text{val}(\mathfrak{G}_{C_m}^{\otimes d}) \leq 1 - (1/m) \cdot \Omega(\sqrt{d/\log d}).$$

Let us make some remarks on this bound, writing $\epsilon = 1/m$ and disregarding the logarithmic factor for clarity. (The reader should thus note that the following remarks only hold up to a logarithmic factor.) First, since $\epsilon \cdot \sqrt{d} = \sqrt{\epsilon^2 \cdot d}$, the former quantity is always larger when the quantities are less than 1; hence our bound improves upon the Feige-Lovász SDP bound (2) in an essential way. (See Figure 3 for a sketch.) It also circumvents the $1 - \epsilon^2 \cdot d$ bound suggested by Feige for Raz/Holenstein-style proofs. Second, our bound has a strong guarantee that is not present in (2): namely, it gives something nontrivial for small, even constant d . By way of contrast, (2) has no content until d is bigger than $1/\epsilon$.

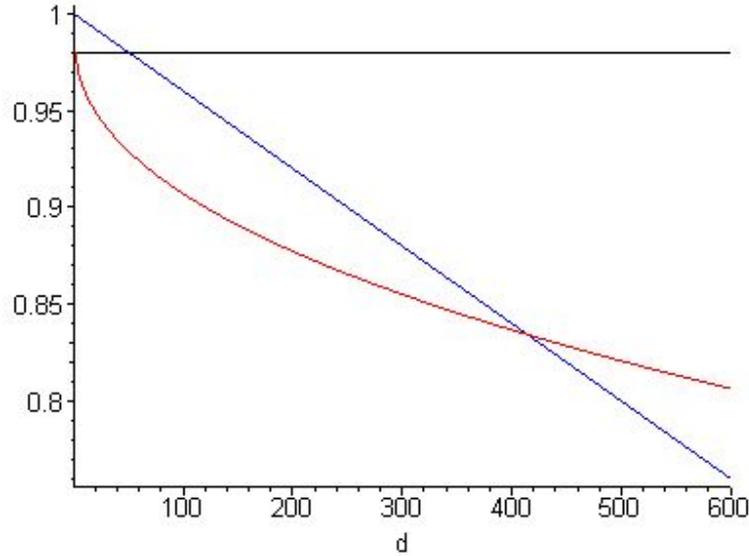


Figure 3: $\text{val}(\mathfrak{G}_{C_{33}}^{\otimes d})$ vs. d . Trivial bound in black, SDP bound in blue, our bound in red.

Finally, however, our bound just barely fails to improve the asymptotic relationship of $1 - \epsilon$ vs. $1 - \Theta(\sqrt{\epsilon})$ gaps, the necessary condition for showing that $1 - \epsilon$ vs. $1 - \Omega(\sqrt{\epsilon})$ NP-hardness for Max-Cut implies the Unique Games Conjecture via parallel repetition. This is because for any number of parallel repetitions d , our theorem transforms a $1 - \epsilon$ vs. $1 - \Theta(\sqrt{\epsilon})$ gap into a $1 - \Theta(\epsilon \cdot d)$ vs. $1 - \Theta(\sqrt{\epsilon \cdot d})$ gap, which merely maintains the square-root gap relationship rather than improves it.

We end this section by showing that determining $\text{val}(\mathfrak{G}_{C_m}^{\otimes d})$ is equivalent to determining the solution $\delta'(d, m)$ to the Odd Cycle Elimination Problem on $K_2 \times (\mathbb{Z}_m^d)_\infty$:

Proposition 2.4 $\text{val}(\mathfrak{G}_{C_m}^{\otimes d}) = 1 - \delta'(d, m)$.

Proof: Say that a *labeling* of $K_2 \times (\mathbb{Z}_m^d)_\infty$ is a function f mapping the vertices of $K_2 \times (\mathbb{Z}_m^d)_\infty$ into \mathbb{Z}_2^d . Given an edge in this graph connecting (A, x) to $(B, x + \lambda)$ (where $x \in \mathbb{Z}_m^d$, $\lambda \in \{-1, 0, 1\}^d$), say that it is *satisfied* by the labeling if $f(A, x) - f(B, x + \lambda) = \lambda \pmod{2}$ and is *violated* otherwise. We now have that the quantity $1 - \text{val}(\mathfrak{G}_{C_m}^{\otimes d})$ is equal to the minimum over all labelings of the fraction of violated edges.

If \mathcal{E} is a set of edges of $K_2 \times (\mathbb{Z}_m^d)_\infty$, say that it is a *blocking solution* if deletion of the edges \mathcal{E} yields a graph $(K_2 \times (\mathbb{Z}_m^d)_\infty) \setminus \mathcal{E}$ with no topologically odd cycles. The proposition follows from showing two things: (i) in any labeling of $K_2 \times (\mathbb{Z}_m^d)_\infty$, the set of violated edges form a blocking solution; and, (ii) if \mathcal{E} is a blocking solution then there is a labeling of $K_2 \times (\mathbb{Z}_m^d)_\infty$ so that the only violated edges are in \mathcal{E} .

For (i), given any labeling f , let \mathcal{E} be the set of violated edges. Suppose by way of contradiction that \mathcal{E} is not a blocking solution. Then there must be some cycle L in $(K_2 \times G) \setminus \mathcal{E}$ and a coordinate $1 \leq i \leq d$

such that $[L]_i$ is odd. Projecting L into $(\mathbb{Z}_m^d)_\infty$ and restricting attention to the cycle's i th coordinate, we get a cycle in the graph $(\mathbb{Z}_m^1)_\infty$ which “wraps around” an odd number of times. It's easy to see that in going along this cycle, the number of length-1 steps (i.e., the number of non-self-loop steps) must be odd. Now consider the sequence of i th coordinates of the labels f gives to the vertices along L . Since all edges in L are satisfied, this sequence must change value (in \mathbb{Z}_2) an odd number of times. But this is impossible since it starts and ends at the same value.

For (ii), given a blocking solution \mathcal{E} , delete its edges from $K_2 \times (\mathbb{Z}_m^d)_\infty$; we claim it is now possible to label the vertices so that all remaining edges are satisfied. In particular, label any starting vertex arbitrarily, and then “greedily” propagate labels along the edges remaining in $(K_2 \times G) \setminus \mathcal{E}$ (all choices are forced). Do this for all connected components of $(K_2 \times G) \setminus \mathcal{E}$. The only way we could get “stuck” is if there exists a cycle L in the graph which can't be labeled in this fashion. But it is easy to see that such a labeling scheme fails for L only if there is a coordinate i such that L makes an odd number of length-1 steps in the i th coordinate. But again, this implies that $[L]_i$ is odd, contradicting the fact that \mathcal{E} is a blocking solution. \square

Remark 2.5 *The Odd Cycle Elimination Problem on $K_2 \times (\mathbb{Z}_m^d)_\infty$ is related to the Cycle Elimination Problem on $(\mathbb{Z}_m^d)_\infty$ as follows: $\delta'(d, m) \leq 2\delta(d, m)$. To see this, let \mathcal{E} be a set of edges in $(\mathbb{Z}_m^d)_\infty$ whose deletion eliminates all topologically nontrivial cycles in $(\mathbb{Z}_m^d)_\infty$, and take*

$$\begin{aligned} \mathcal{E}' = & \{(A, x)-(B, y) : (x, y) \in \mathcal{E}\} \\ & \cup \{(B, x)-(A, y) : (x, y) \in \mathcal{E}\} \end{aligned}$$

to be a set of edges in $K_2 \times (\mathbb{Z}_m^d)_\infty$. Then $|\mathcal{E}'| = 2|\mathcal{E}|$, and clearly deleting \mathcal{E}' from $K_2 \times (\mathbb{Z}_m^d)_\infty$ eliminates all topologically nontrivial cycles (and hence all topologically odd cycles) in $K_2 \times (\mathbb{Z}_m^d)_\infty$.

3 Foams

The problem of determining $A(d)$, the least surface area of a cell that tiles \mathbb{R}^d by the lattice \mathbb{Z}^d , is an example of a *foam* problem. The most famous of all foam problems is the Kelvin Problem, which asks for the foam of minimal surface area foam which partitions \mathbb{R}^3 into cells of volume 1.¹ This problem has been studied since at least 1781 [Lhu81]; it is named after Thomson, a.k.a. Lord Kelvin, who in 1887 [Tho87] proposed the following solution: Take the Voronoi cells of the root lattice A_3^* (appropriately scaled), and then slightly “relax” them so that they satisfy the Plateau rules [Pla73] for soap films. This was believed to be minimal for about a century, until the Weaire-Phelan foam was discovered in 1993 [WP94]. Two centuries ago the Kelvin Problem was described as “one of the most difficult in geometry” [Lhu81] and Morgan has recently speculated (see [Hal00]) that it might take another century to solve. The 2-dimensional version was solved only in 1999, by Hales [Hal01], with the regular hexagon tiling — i.e., the Voronoi cells of the lattice A_2^* — being minimal, as expected. Very little is known about the problem in dimensions $d > 3$.

Our Foam Problem on $\mathbb{R}^d/\mathbb{Z}^d$ is a special case of the d -dimensional Kelvin Problem in which the foam is required to be periodic with period \mathbb{Z}^d . It is also natural to consider the version in which the foam is only required to be periodic with respect to *some* lattice (as the Kelvin foam is, but the Weaire-Phelan foam isn't). For reasons to be described shortly, we call this the *Choe Problem for Flat Tori*.

Choe Problem for Flat Tori: What is the least surface area of a cell that tiles \mathbb{R}^d by a determinant-1 lattice Λ ?

One may restate this problem as, “What is the least surface area of a fundamental domain for a unit-volume flat torus, \mathbb{R}^d/Λ ?” Choe [Cho89] studied this problem for all 3-dimensional smooth Riemannian manifolds, not just flat tori, and for each showed existence of a minimizer. It has been conjectured by Sullivan and others [SM96] that the Kelvin foam solves the Choe Problem for Flat Tori in three dimensions (with Λ being A_3^* , appropriately scaled). Choe's work also completely solves the problem in 2 dimensions; it implies that for any lattice, the optimal cell is the unique hexagon with straight sides and 120° angles; and,

¹In fact, it is not currently known if a minimal such surface exists.

the optimal lattice is the triangular one, leading to the regular hexagon. For $d > 3$ it is not even known if a least-area cell exists; one of the few things known is that in any lattice tiling the cell surfaces are minimal surfaces (i.e., have mean curvature 0) [WKC⁺04]. Choe explicitly posed our Foam Problem on $\mathbb{R}^3/\mathbb{Z}^d$ as an open problem; as far as we know, there is not even a conjecture for the minimizer in this case.

3.1 Cycle Elimination on $(\mathbb{Z}_m^d)_\infty$ reduces to the Foam Problem on $\mathbb{R}^d/\mathbb{Z}^d$

In this section we show that good solutions to the Foam Problem on $\mathbb{R}^d/\mathbb{Z}^d$ translate into good solutions for the Cycle Elimination Problem on $(\mathbb{Z}_m^d)_\infty$, and thus upper bounds for $\text{val}(\mathfrak{G}_{C_m}^{\otimes d})$. A somewhat related relation was observed by Bollobás, Kindler, Leader, and O’Donnell in [BKLO06].

Theorem 3.1 *Let C be a cell with piecewise smooth surface that tiles \mathbb{R}^d by \mathbb{Z}^d and has surface area A . Then $\delta(d, m) \leq 2\sqrt{2/3} \cdot A/m$.*

Proof: Let \mathcal{S} denote the boundary surface of C modulo \mathbb{Z}^d . Note that \mathcal{S} has area $A/2$, since the boundary of C overlaps with itself when it is placed inside the torus. Further, note that \mathcal{S} has the property that every homotopically nontrivial loop inside \mathcal{T}^d touches \mathcal{S} ; for otherwise, it could be pulled back to an unbounded path in \mathbb{R}^d avoiding the foam generated by C ’s boundary.

Let $T \in [0, 1)^d$ be a random translate in \mathcal{T}^d , and define \mathcal{S}' to be $T + \mathcal{S}$; i.e., \mathcal{S} translated by T within \mathcal{T}^d . Certainly \mathcal{S}' also has the property that every homotopically nontrivial loop in \mathcal{T}^d touches it. Now embed the graph $(\mathbb{Z}_m^d)_\infty$ in \mathcal{T}^d in the natural way, with vertices as points and its edges as straight line segments. Finally, let \mathcal{E} denote the set of all edges that intersect \mathcal{S}' (a random set depending on T). It’s clear that deleting the edges \mathcal{E} from the graph $(\mathbb{Z}_m^d)_\infty$ eliminates all topologically nontrivial cycles in $(\mathbb{Z}_m^d)_\infty$, since cycles in $(\mathbb{Z}_m^d)_\infty$ correspond to actual piecewise straight loops in \mathcal{T}^d with the same homotopy class. Hence we may complete the proof by showing

$$\mathbf{E}_T [|\mathcal{E}| / ((3m)^d / 2)] \leq 2\sqrt{2/3} \cdot A$$

(since $(3m)^d / 2$ is the number of edges in $(\mathbb{Z}_m^d)_\infty$), or equivalently,

$$\mathbf{E}_T [|\mathcal{E}|] \leq 2\sqrt{2/3} \cdot 3^d \cdot \text{area}(\mathcal{S}). \quad (3)$$

To do this we decompose \mathcal{S} into finitely many small patches. Specifically, we write $\mathcal{S} = \bigcup_{i=1}^N \mathcal{P}_i$, where:

1. \mathcal{P}_i is a smooth closed oriented surface with piecewise smooth boundary;
2. for each i , the Euclidean distance between every pair of points in \mathcal{P}_i is at most $1/2$;
3. $\sum_{i=1}^N \text{area}(\mathcal{P}_i) = \text{area}(\mathcal{S})$.

Now to prove (3), by linearity of expectation it suffices to show that

$$\mathbf{E}_T [\# \text{ edges that touch } T + \mathcal{P}_i] \leq 2\sqrt{2/3} \cdot 3^d \cdot \text{area}(\mathcal{P}_i) \quad (4)$$

for all i . Let $\lambda \in \{-1, 0, 1\}^d$ be an edge “type”. By property (2) above, any translate $T + \mathcal{P}_i$ touches at most two edges of type λ (since given three points on distinct edges of the same type, some two must be at distance at least 1). It follows that

$$\mathbf{E}_T [\# \text{ edges that touch } T + \mathcal{P}_i] \leq 2 \sum_{\lambda \in \{-1, 0, 1\}^d} \mathbf{Pr}_T [T + \mathcal{P}_i \text{ touches an edge of type } \lambda].$$

Further, $T + \mathcal{P}_i$ touches at least one edge of type λ if and only if $T + (\mathcal{P}_i \oplus \lambda)$ contains a lattice point, where $(\mathcal{P}_i \oplus \lambda)$ denotes the Minkowski sum of \mathcal{P}_i and the line segment joining 0 and λ . The probability of this is

$$\text{area}(\mathcal{P}_i \oplus \lambda) \leq \int_{\mathcal{P}_i} |\vec{n} \cdot \lambda| dS,$$

where \vec{n} denotes the unit normal to the surface. We conclude

$$\begin{aligned} \mathbf{E}_T [\# \text{ edges that touch } T + \mathcal{P}_i] &\leq 2 \int_{\mathcal{P}_i} \left(\sum_{\lambda} |\vec{n} \cdot \lambda| \right) dS \\ &\leq 2 \cdot 3^d \cdot \int_{\mathcal{P}_i} \sqrt{\mathbf{E}_{\lambda} [(\vec{n} \cdot \lambda)^2]} dS = 2 \cdot 3^d \cdot \int_{\mathcal{P}_i} \sqrt{2/3} dS = 2\sqrt{2/3} \cdot 3^d \cdot \text{area}(\mathcal{P}_i), \end{aligned}$$

where the second inequality is by Cauchy-Schwarz and the following equality is by direct calculation, using the fact that \vec{n} is a unit vector. This is the desired (4), completing the proof. \square

3.2 Bounds on $A(d)$

In this section we give bounds on $A(d)/2$, where $A(d)$ denotes the least surface area of a cell that tiles \mathbb{R}^d by \mathbb{Z}^d .

When $d = 2$, the work of Choe [Cho89] implies that the optimal cell is the hexagon with 120° angles pictured in Figure 4. This implies $A(2)/2 = (1 + \sqrt{3})/\sqrt{2} \approx 1.93$.

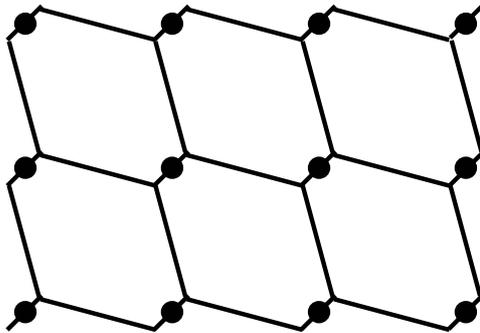


Figure 4: Solution to the Foam Problem on $\mathbb{R}^2/\mathbb{Z}^2$

For general d , we essentially only know the “obvious” bounds. For an upper bound we have $A(d)/2 \leq d$, since the unit cube tiles \mathbb{R}^d by \mathbb{Z}^d and has surface area $2d$. It is possible to improve this very slightly to $d - \exp(-O(d \log d))$ by inserting a radius- $1/2$ sphere at the corners of the cubic tiling (the details will appear in the full version of this paper). Although this gives no significant improvement on the bound, it does show that the cube is not optimal.

For a lower bound, the isoperimetric inequality implies that *any* cell of volume 1 has surface area at least that of the ball of volume 1. An easy calculation using Stirling’s formula shows that in d dimensions the ball of unit volume has surface area $\sqrt{2\pi e} \sqrt{d} (1 \pm o(1))$. This gives the lower bound $A(d)/2 \geq \sqrt{\pi e/2} \sqrt{d} - o(\sqrt{d})$.

Another way to prove an $\Omega(\sqrt{d})$ lower bound is as follows: Suppose C tiles $\mathbb{R}^d/\mathbb{Z}^d$. As in the proof of Theorem 3.1, place C inside the unit cubic torus $(\mathbb{Z}_m^d)_\infty^d$ and let \mathcal{S} denote its boundary surface therein. Again, \mathcal{S} has the property that it blocks all topologically nontrivial loops, and so in particular it blocks all of those loops that simply travel straight, parallel to a coordinate axis. As a consequence we have that when \mathcal{S} is projected onto any one of the d lower-dimensional faces of $[0, 1]^d$, the projection must be onto. Hence if one integrates $|\vec{n} \cdot \vec{e}_i|$ along the surface of \mathcal{S} (with \vec{n} denoting a unit normal and \vec{e}_i denoting $(0, \dots, 0, 1, 0, \dots, 0)$), the integral must be at least 1, for each i . Summing this over $i = 1 \dots d$ yields that the integral of $\|\vec{n}\|_1$ over \mathcal{S} is at least d . But since $\|\vec{n}\|_1 \leq \sqrt{d}$, one concludes that the area of \mathcal{S} must be at least $d/\sqrt{d} = \sqrt{d}$. This

proves $A(d)/2 \geq \sqrt{d}$.

As remarked, this lower bound is in fact worse than the isoperimetric inequality lower bound $A(d)/2 \geq \sqrt{\pi e/2} \sqrt{d}(1 - o(1)) \approx 2.066\sqrt{d}$. But this is the method that we manage to translate back to a lower bound for the Cycle Elimination Problem on $(\mathbb{Z}_m^d)_\infty$, which is discussed in the next section.

Regarding the generalized problem in d dimensions, i.e. Choe's Problem for Flat Tori, Sullivan remarked to us that a natural candidate for a good solution would be the Voronoi cell for the lattice A_d^* (appropriately scaled); the reason is that this yields the best solution for $d = 2, 3$. In Appendix A we show that this cell has surface area $2K \cdot d \pm O(d^{2/3})$, where

$$K = \sum_{R=1}^{\infty} \frac{R^{R-3/2}}{R!} e^{-R} \approx .6186.$$

Thus this cell is only marginally better than the unit cube.

Finally, we would like to remark that for the Foam Problem for $\mathbb{R}^3/\mathbb{Z}^3$, which Choe left as an open problem, we know no bounds better than $(9\pi/2)^{1/3} \leq A(3)/2 \leq 3$, and know no conjectured best solution.

4 A lower bound for the Cycle Elimination Problem

The remainder of this paper is devoted to proving our main theorem, a lower bound for the Cycle Elimination Problem on $(\mathbb{Z}_m^d)_\infty$ (and for the Odd Cycle Elimination Problem on $K_2 \times (\mathbb{Z}_m^d)_\infty$):

Theorem 4.1 *Provided $m \geq \sqrt{d/\log d}$,*

$$\delta(d, m) \geq (1/m) \cdot \Omega(\sqrt{d/\log d}), \tag{5}$$

and indeed $\delta'(d, m) \geq (1/m) \cdot \Omega(\sqrt{d/\log d})$ as well.

We concentrate for now on proving (5) — as observed in Claim 4.8, our proof technique actually gives the lower bound on $\delta'(d, m)$ for free.

Idea of the proof: Our strategy for proving Theorem 4.1 will be to try to “discretify” the \sqrt{d} lower bound for the Foam Problem on $\mathbb{R}^d/\mathbb{Z}^d$ that comes from showing that all straight axis-parallel loops must be blocked. This discretization turns out to be quite elaborate. Obviously it no longer suffices just to look at axis-parallel loops — i.e., loops formed only by using edges of “type” $\lambda = (0, \dots, 0, 1, 0, \dots, 0)$ — since these constitute only a $1/3^d$ or so fraction of all edges. Instead, we look at “essentially” axis-parallel loops, i.e. cycles that stay within narrow axis-parallel tubes.

Let us introduce some definitions that were first considered in the work of Saks, Samorodnitsky and Zosin [SSZ04] (and later in [BKLO06]) on the *vertex-deletion* version of the Cycle Elimination Problem on $(\mathbb{Z}_m^d)_\infty$. In the following definition, and in the rest of the paper, we call elements of \mathbb{Z}_m^d points *vertices* and *points*.

Definition 4.2 *Let $1 \leq i \leq d$ be a direction. We introduce the following kinds of subsets of vertices of $(\mathbb{Z}_m^d)_\infty$:*

- *A section in direction i : a set of the form $x + \{0, 1\}^{i-1} \times \{0\} \times \{0, 1\}^{d-i}$ for some $x \in \mathbb{Z}_m^d$. The point x is called the base point of the section.*
Two sections are called successive if they are both in the same direction i and the difference between their base points is $(0, \dots, 0, 1, 0, \dots, 0)$.
- *A tube in direction i : a set of the form $x + \{0, 1\}^{i-1} \times \mathbb{Z}_m \times \{0, 1\}^{d-i}$ for some $x \in \mathbb{Z}_m^d$.*
(Note that a tube T is the union of m successive sections oriented in the same direction as T .) We say that a section belongs to a tube T , or that it is contained in T , if it is contained in T and has the same direction as T .

- A cube: a set of the form $x + \{0, 1\}^d$ for some $x \in \mathbb{Z}_m^d$.

Note that a cube can be obtained in d many ways as a union of two successive sections.

We have that a section contains 2^{d-1} points, a cube contains 2^d points, and a tube contains $m2^{d-1}$ points. The reader may find it helpful to visualize these definition in 3 dimensions (and think of a section as a “square”).

Fix \mathcal{E} to be a set of “deleted edges” that block all topologically nontrivial cycles in $(\mathbb{Z}_m^d)_\infty$. Our goal is to show that the fractional size of \mathcal{E} , i.e. $|\mathcal{E}|/((3m)^d/2)$, is bounded below by $(1/m) \cdot \Omega(\sqrt{d/\log d})$. We do this by considering the vertices which have a constant fraction of their edges in \mathcal{E} , and showing that they constitute a $(1/m) \cdot \Omega(\sqrt{d/\log d})$ fraction of all vertices.

Definition 4.3 A vertex x in $(\mathbb{Z}_m^d)_\infty$ is said to be marked if at least a 10^{-6} fraction of the edges touching it are in \mathcal{E} .

The following lemma immediately implies (5).

Lemma 4.4 The fraction of vertices in $(\mathbb{Z}_m^d)_\infty$ that are marked is at least $(1/m) \cdot \Omega(\sqrt{d/\log d})$.

The basis of our proof of Lemma 4.4, and thus of (5) is the following: If we are to eliminate all topologically nontrivial cycles in $(\mathbb{Z}_m^d)_\infty$, then at the very least, for each tube T we must eliminate all cycles that “wrap around once along the vertices of $(\mathbb{Z}_m^d)_\infty$ ”.

The proof of Lemma 4.4 now divides into two parts. First we argue in Theorem 4.6 that *every* tube T must contain some section where the fraction of marked points is very large — at least 95%.² Otherwise, we show that one can find cycles that march through the successive sections of T . In the second part, Theorem 4.7, we show that if every tube contains a section with 95% of its points marked then the overall fraction of marked vertices in $(\mathbb{Z}_m^d)_\infty$ is at least $(1/m) \cdot \Omega(\sqrt{d/\log d})$ as required for Lemma 4.4. In order to formalize the statements in these two parts, we need to define what we mean by cycles that “wrap around once along the vertices of $(\mathbb{Z}_m^d)_\infty$ ”.

Definition 4.5 Let L be a cycle in $(\mathbb{Z}_m^d)_\infty$, and T be a tube in it. Write the successive sections of T as $S_0, S_1, \dots, S_{m-1}, S_m = S_0$. We say that L is an even T -cycle if it is a concatenation of paths $P_0, Q_0, P_1, Q_1, \dots, P_{m-1}, Q_{m-1}$ with the following properties:

1. For every i , P_i is the concatenation of length-2 paths connecting vertices in S_i .
2. For every i , Q_i is a length-2 path connecting a vertex in S_i to a vertex in S_{i+1} .

It is easy to see that if the tube T is in direction i , then the homotopy class of an even T -cycle is $(0, \dots, 0, 1, 0, \dots, 0)$, where the 1 is in the i th coordinate.

Lemma 4.4, and thus (5), is immediately implied by the following two theorems.

Theorem 4.6 Let \mathcal{E} be a set of deleted edges in $(\mathbb{Z}_m^d)_\infty$ and let T be a tube in $(\mathbb{Z}_m^d)_\infty$ such that every section in T has less than 95% of its points marked. Then there exists an even T -cycle in $(\mathbb{Z}_m^d)_\infty \setminus \mathcal{E}$.

Theorem 4.7 Let m and d be integers such that $m \geq \sqrt{d/\log d}$. Let \mathcal{V} be a set of marked vertices in $(\mathbb{Z}_m^d)_\infty$ such that every tube in $(\mathbb{Z}_m^d)_\infty$ contains a section with at least 95% of its points marked. Then the overall fraction of marked points, $|\mathcal{V}|/m^d$, is at least $(1/m) \cdot \Omega(\sqrt{d/\log d})$.

In fact, the two theorems above also imply the same lower bound on $\delta'(d, m)$, and thus our main theorem, Theorem 4.1.

Claim 4.8 Theorem 4.6 and Theorem 4.7 imply that $\delta'(d, m) \geq (1/m) \cdot \Omega(\sqrt{d/\log d})$.

²Our proof is such that we can make this quantity $1 - \eta$ at the expense of making the “marked” fraction $\Omega(\eta^3)$.

Proof: Let \mathcal{E}' be a set of deleted edges in $K_2 \times (\mathbb{Z}_m^d)_\infty$ of fractional size α , and suppose that \mathcal{E}' blocks all topologically odd cycles. We would like to show that the lower bound in (5) applies to α . Project the edges of \mathcal{E}' into $(\mathbb{Z}_m^d)_\infty$ in the natural way, obtaining a set \mathcal{E} of edges in $(\mathbb{Z}_m^d)_\infty$ of fractional size at most 2α . We argue that \mathcal{E} blocks all even T -cycles in $(\mathbb{Z}_m^d)_\infty$.

Otherwise, there is a tube T in $(\mathbb{Z}_m^d)_\infty$ and an even T -cycle L which is not blocked by \mathcal{E} . Because this cycle is of even length, we can lift it into a cycle L' in $K_2 \times (\mathbb{Z}_m^d)_\infty$, which is clearly not blocked by \mathcal{E}' . Also, being an even T -cycle, the homotopy class of L is a unit vector and so L is topologically odd; thus L' is also topologically odd, since it projects to L . But this contradicts the assumption that \mathcal{E}' blocks all topologically odd cycles in $K_2 \times (\mathbb{Z}_m^d)_\infty$.

Since \mathcal{E} blocks all even T -cycles, Theorem 4.6 and Theorem 4.7 now easily imply that $2\alpha \geq (1/m) \cdot \Omega(\sqrt{d/\log d})$, completing the proof of the claim. \square

In the next two sections we prove Theorem 4.6 and Theorem 4.7, completing the proof of our main theorem, Theorem 4.1.

4.1 Proof of Theorem 4.6

The main tool used to prove this theorem is the recent isoperimetric inequality on the discrete cube due to Mossel, O'Donnell, Regev, Steif and Sudakov [MOR⁺06], based on the reverse Bonami-Beckner inequality.

We show the existence of an even T -cycle L in $(\mathbb{Z}_m^d)_\infty \setminus \mathcal{E}$ in two steps. First we show that every section in T contains a ‘‘giant’’ component of unmarked vertices. More specifically, we show that for every section in T there exists a set of $2\% \cdot 2^{d-1}$ unmarked vertices, every two of which are connected by a path in $(\mathbb{Z}_m^d)_\infty \setminus \mathcal{E}$ that has the P_i type described in Definition 4.5. We then show that the giant components in every pair of successive sections are connected by paths of length 2. Using the paths within the giant components and the length-2 paths between them, one easily constructs an even T -cycle.

Giant components: To accomplish the first step, we show that in a particular section, if A and B are any two sets of unmarked vertices each of size at least $2\% \cdot 2^{d-1}$, then there is a path of length 2 between A and B . This implies the existence of the desired connected component of unmarked vertices of size $(1 - 95\%) - 2\% > 2\%$. So suppose that T is in the direction i . Now any unmarked point has only a 10^{-6} fraction of its edges in \mathcal{E} , so at most a $3 \cdot 10^{-6}$ fraction of its edges *with 0 displacement in the i th direction* are in \mathcal{E} . So we can reduce to a problem on the $(d - 1)$ -dimensional lattice, and in particular, the following claim establishes the existence of the desired giant components (when applied with $c = d - 1$):

Claim 4.9 *Let $A, B \subseteq \{0, 1\}^c \subseteq \mathbb{Z}_m^c$ be subsets of size at least $2\% \cdot 2^c$ each, and suppose that a set \mathcal{E} of deleted edges contains at most $3 \cdot 10^{-6}$ of the edges of each vertex in $A \cup B$. Then there exists a path of length 2 from A to B .*

Proof: Pick a random length-2 path by uniformly choosing a random point $x \in \{0, 1\}^c$, then walking along a random edge from that point into $\{-1, 0, 1, 2\}^c$ (we allow walking on deleted edges), and then taking a uniformly random edge from that point back into a point $y \in \{0, 1\}^c$. Note that this yields the uniform probability distribution on y and that the second edge in this path is uniformly distributed over all edges touching y .

We show that the probability that x is in A and that y is in B is at least $(2\%)^3 > 6 \cdot 10^{-6}$. Then the union bound implies that with positive probability, the random length-2 path connects A and B while avoiding \mathcal{E} , thereby completing the proof of Claim 4.9.

What is the actual distribution on (x, y) ? The point x is chosen uniformly, and y can be chosen by holding each bit of x steady with probability $1/3$ and re-randomizing it with probability $2/3$. Thus the distribution on (x, y) is that of a ‘‘(1/3)-correlated pair’’, in the notation of [MOR⁺06]. We now use the following instantiation of the isoperimetric inequality in [MOR⁺06, Corollary 3.5] to lower-bound the probability that $x \in A$ and $y \in B$:

Theorem 4.10 *Let $A, B \subseteq \{0, 1\}^c$ be sets of fractional size σ . Let $x, y \in \{0, 1\}^c$ be a ρ -correlated pair. Then*

$$\Pr[x \in A, y \in B] \geq \sigma^{2/(1-\rho)}.$$

Applying this result in our situation, where $\sigma = 2\%$, $\rho = 1/3$, and $2/(1-\rho) = 3$, yields the lower bound of $(2\%)^3 > 6 \cdot 10^{-6}$ as claimed. \square

Connecting giant components: We now want to show that there is a length-2 path between the giant components in every two successive sections. The proof of this is extremely similar to the one we used to establish the existence of the giant components; we now let A and B be the giant components in the successive sections. The only thing we need to do differently is observe that every point in A has at most a $3 \cdot 10^{-6}$ fraction of its edges *that have displacement +1 in the i th direction* not in \mathcal{E} . We now pick our random length-2 walks by choosing x uniformly at random from A 's section, picking a uniformly random edge from x that has displacement +1 in the i th direction, and then walking along a uniformly random edge into B 's section.

The proof of Theorem 4.6 is complete. \square

4.2 Proof of Theorem 4.7

In this section we prove Theorem 4.7, but first, let us note that it is essentially tight.

Tightness of Theorem 4.7: Suppose we mark all points in $(\mathbb{Z}_m^d)_\infty$ whose coordinates sum, modulo m , to a number in the range $[d/2 - C\sqrt{d}, d/2 + C\sqrt{d}]$, for some large constant C . This clearly marks only a $(1/m) \cdot O(\sqrt{d})$ fraction of the points. But now in each tube, if we consider the sections whose base point x has coordinate-sum $0 \pmod{m}$, they all will have an extremely large fraction of their points marked. (Indeed, we can make the fraction more than 95% by taking $C = 1$.)

To prove Theorem 4.7, assume that we have a set of marked points that satisfies the hypothesis of the theorem. We now choose *exactly* one section per tube with at least a 95% fraction of marked points; call these *marked sections*. (There may now be unmarked sections with more than a 95% fraction of their points marked.) Write \mathcal{S} for this collection of marked sections. The idea behind the proof is to show that these sections cannot overlap too much.

Definition 4.11 *Given a collection \mathcal{S} of marked sections in $(\mathbb{Z}_m^d)_\infty$, define the degree of a point $x \in \mathbb{Z}_m^d$, denoted $\deg(x)$, to be the number of marked sections in which it participates, $\#\{S \in \mathcal{S} : x \in S\}$.*

The bulk of our work will go into proving the following:

Lemma 4.12 *Let \mathcal{S} be a collection of marked sections, exactly one section per tube of $(\mathbb{Z}_m^d)_\infty$, and let $4/m \leq \tau \leq 1$. Then the number of points in \mathbb{Z}_m^d with degree at least $\tau 2^{d-1}$ is at most $\text{poly}(1/\tau) \cdot \exp(-\Omega(\tau^2 d)) \cdot m^{d-1}$.*

Let us see why Lemma 4.12 implies Theorem 4.7:

Proof: (of Theorem 4.7) Let \mathcal{S} be a collection of exactly one section per tube, and assume that for every section $S \in \mathcal{S}$ at least a 95% fraction of its 2^{d-1} points are marked. Since each marked point has degree at least 1, we conclude that

$$\sum_{x \in (\mathbb{Z}_m^d)_\infty} \deg(x) \geq |\mathcal{S}| \cdot 95\% \cdot 2^{d-1} = 95\% \cdot d \cdot (2m)^{d-1}.$$

Now apply Lemma 4.12 with $\tau = C\sqrt{(\log d)/d}$, where C is chosen large enough so that $\text{poly}(1/\tau) \cdot \exp(-\Omega(\tau^2 d)) < 1/d^2$ (note that indeed $\tau \geq 4/m$, by our assumption on m). We conclude that the number

of points with degree at least $D = \tau d 2^{d-1} = O(\sqrt{d \log d}) 2^{d-1}$ is at most $(1/d^2)m^{d-1}$. Since the maximum possible degree is $d 2^{d-1}$, these points contribute at most $(1/d)(2m)^{d-1}$ to $\sum_x \deg(x)$. Thus

$$\begin{aligned} \sum_{x : \deg(x) \leq D} \deg(x) &\geq (95\% \cdot d - 1/d) \cdot (2m)^{d-1} \\ &= \Omega(d) \cdot (2m)^{d-1}, \end{aligned}$$

which implies that the number of points x with nonzero degree is at least $(1/D)\Omega(d) \cdot (2m)^{d-1} = \Omega(\sqrt{d/\log d})m^{d-1}$, as claimed. \square

The remainder of this section is devoted to the proof of Lemma 4.12.

Assume \mathcal{S} is a collection of sections consisting of exactly one section per tube of $(\mathbb{Z}_m^d)_\infty$, and let \mathcal{V} denote the set of points in \mathbb{Z}_m^d that have degree at least $\tau d 2^{d-1}$ with respect to \mathcal{S} . Our goal is to upper-bound the cardinality of \mathcal{V} .

We begin with some definitions:

Definition 4.13 *The section distribution, denoted $\mathcal{D}_{\text{section}}$, is the distribution on sections gotten by first drawing x uniformly from \mathcal{V} and then choosing a uniformly random section containing x . We define similarly the cube distribution $\mathcal{D}_{\text{cube}}$ and the tube distribution $\mathcal{D}_{\text{tube}}$.*

Remark 4.14 *Let T be drawn from $\mathcal{D}_{\text{tube}}$ and let x be a uniformly random point in $\mathcal{V} \cap T$ (which is a nonempty set). Then the distribution on (x, T) is the same as if we had drawn x uniformly from \mathcal{V} and let T be a random tube containing x . This uses the fact that each point in \mathcal{V} is contained in the same number of tubes.*

Definition 4.15 *The \mathcal{V} -density of a section S is $|\mathcal{V} \cap S|/|S| = |\mathcal{V} \cap S|/2^{d-1}$. We define analogously the \mathcal{V} -density of cubes and tubes.*

Definition 4.16 *The relative section density sequence of a tube T is the circular sequence (p_0, \dots, p_{m-1}) in which $p_i = |\mathcal{V} \cap T_i|/|\mathcal{V} \cap T|$, where T_i denotes the i th section of T . The maximum relative section density of T , denoted $\rho(T)$, is $\max_i p_i$.*

The first step in proving Lemma 4.12 is showing that a typical tube T from $\mathcal{D}_{\text{tube}}$ has its \mathcal{V} points somewhat concentrated on one section — i.e., has large $\rho(T)$. The intuitive reason for this is that if T has its \mathcal{V} points spread evenly throughout its sections, then most of them will not be in T 's marked section, contradicting the idea that the points of \mathcal{V} have high degree. Formally:

Proposition 4.17 $\mathbf{E}_{T \leftarrow \mathcal{D}_{\text{tube}}} [P(T)] \geq \rho$.

Proof: Pick T according to $\mathcal{D}_{\text{tube}}$ and then pick x to be a random point from $\mathcal{V} \cap T$. Say the pair (x, T) is “successful” if x is in T 's marked section from \mathcal{S} . Conditioned on T , the probability of a successful draw is at most $\rho(T)$; hence the overall probability of a successful draw is at most $\mathbf{E}_{T \leftarrow \mathcal{D}_{\text{tube}}} [\rho(T)]$.

On the other hand, by Remark 4.14 we can view T as being a uniformly random tube containing x . But by definition of $x \in \mathcal{V}$, the probability a random tube through x has x in its marked section is at least ρ . Hence the overall probability of a successful draw must be at least ρ . This completes the proof. \square

Next, we show that if a tube's maximum relative section density is large, then it must have two successive sections in which the first section has a much larger \mathcal{V} -density than the second.

Definition 4.18 *Given two densities $p, q \in [0, 1]$, define their imbalance to be*

$$\text{imb}(p, q) = 1 - \frac{\min\{p, q\}}{\max\{p, q\}}.$$

WE have $0 \leq \text{imb}(p, q) \leq 1$, with $\text{imb}(p, q) = 1$ if $p = q$ and $\text{imb}(p, q) = 0$ if one of p or q is zero and the other is nonzero. We leave $\text{imb}(0, 0)$ undefined.

Proposition 4.19 *Let (p_0, \dots, p_{m-1}) be a circular sequence of numbers in $[0, 1]$ summing to 1. Let $\rho = \max_i p_i$ and assume $m \geq 2/\rho$. Suppose we pick $0 \leq J < m$ with probability p_J . Then $\mathbf{E}[\text{imb}(p_J, p_{J+1})^2] \geq \Omega(\rho^2)$.³*

Proof: Assume without loss of generality that $p_0 = \rho$. Let k be as small as possible such that $p_k \leq \rho/2$; such a k must exist and be at most $2/\rho$, since $m \geq 2/\rho$ and the p_i 's sum to 1. If there is some $0 \leq i < k$ such that $\text{imb}(p_i, p_{i+1}) \geq 0.1$ then we are certainly done (with the better bound of $\Omega(\rho)$). Otherwise each ratio p_i/p_{i+1} is in the range $[0.9, 1/0.9]$ which means that $\text{imb}(p_i, p_{i+1}) = \Theta(|\log(p_i/p_{i+1})|)$. Now

$$\begin{aligned} \sum_{i=0}^{k-1} \text{imb}(p_i, p_{i+1})^2 &= \Theta(1) \cdot \sum_{i=0}^{k-1} \log^2(p_i/p_{i+1}) \\ &\geq \Theta(1/(k-1)) \left(\sum_{i=0}^{k-1} \log(p_i/p_{i+1}) \right)^2 && \text{(Cauchy-Schwarz)} \\ &= \Theta(1/k) \log^2(p_0/p_k) \geq \Omega(1/k) = \Omega(\rho). \end{aligned}$$

Since each $0 \leq i < k$ is chosen to be J with probability $p_i > \rho/2$, the proof is complete. \square

Combining Propositions 4.17 and 4.19 lets us conclude that the distribution of \mathcal{V} points in a typical cube from the cube distribution is very lopsided.

Definition 4.20 *Given a cube C , its total squared imbalance is $\text{tsi}(C) = \sum_{i=1}^d \text{imb}(p_i, 1 - p_i)^2$, where $p_i = |\mathcal{V} \cap C_i|/|\mathcal{V} \cap C|$ and C_i is either of the two direction- i sections comprising C .*

Proposition 4.21 $\mathbf{E}_{C \leftarrow \mathcal{D}_{\text{cube}}}[\text{tsi}(C)] \geq \Omega(\tau^2 d)$.

Proof: Let x be a random point in \mathcal{V} and T a random tube containing x . Let (p_0, \dots, p_{m-1}) be T 's circular sequence of relative edge densities and let $\rho = \rho(T) = \max_i p_i$. Let I denote the index of x 's section of T . By linearity of expectation, the fact that a random tube is equally likely to be in each of the d directions, and the fact that $\text{imb}(\cdot, \cdot)$ is scale-invariant (i.e., it only depends on the ratio of its arguments), it suffices for us to show that $\mathbf{E}[\text{imb}(p_I, p_{I+1})^2] \geq \Omega(\tau^2)$.

By Proposition 4.19,

$$\begin{aligned} \mathbf{E}[\text{imb}(p_I, p_{I+1})^2 \mid \rho] &\geq \begin{cases} \Omega(\rho^2) & \text{if } \rho \geq 2/m, \\ 0 & \text{else} \end{cases} \\ &= \Omega(1) \cdot (\rho \cdot \mathbf{1}_{[\rho \geq 2/m]})^2. \end{aligned}$$

Now take expectation over ρ and use Cauchy-Schwarz to get

$$\begin{aligned} \mathbf{E}[\text{imb}(p_I, p_{I+1})^2] &\geq \\ &\geq \Omega(1) \cdot \mathbf{E}[\rho \cdot \mathbf{1}_{[\rho \geq 2/m]}]^2 \\ &= \Omega(1) \cdot (\mathbf{E}[\rho] - \mathbf{E}[\rho \cdot \mathbf{1}_{[\rho < 2/m]}])^2 \\ &\geq \Omega(1) \cdot (\tau - 2/m)^2 && \text{(Proposition 4.17)} \\ &\geq \Omega(1) \cdot (\tau - \tau/2)^2 && \text{(since } \tau \geq 4/m). \end{aligned}$$

Thus $\mathbf{E}[\text{imb}(p_I, p_{I+1})^2] \geq \Omega(\tau^2)$, as needed. \square

We now prove that a cube with high tsi contains very few points from \mathcal{V} . This follows from an argument of Fourier analysis of Boolean functions that may be of independent interest. First, we need a lemma of Talagrand [Tal96, Prop 2.2]:

³Note that by definition of I we need not worry about $\text{imb}(0, 0)$ here.

Lemma 4.22 (Talagrand) *Let $f : \{0, 1\}^d \rightarrow \{0, 1\}$ satisfy $\mathbf{E}[f] = \mu$. Then*

$$\sum_{i=1}^d \hat{f}(\{i\})^2 \leq O(\mu^2 \log(e/\mu)).$$

We use this to prove:

Theorem 4.23 *Let C be a d -dimensional discrete cube having $\text{tsi}(C)$ with respect to a set of points \mathcal{V} . Then $|\mathcal{V} \cap C|/|C| \leq \exp(-\Omega(\text{tsi}(C)))$.*

Proof: Let $f : \{0, 1\}^d \rightarrow \{0, 1\}$ denote the indicator function of $\mathcal{V} \cap C$. We need to prove that $\mathbf{E}[f] \leq \exp(-\Omega(\text{tsi}(C)))$. Writing p_i as in the definition of tsi , we have by definition

$$|\hat{f}(\{i\})| = |p_i\mu - (1 - p_i)\mu| = \mu \cdot \max\{p_i, 1 - p_i\} \cdot \text{imb}(p_i, 1 - p_i) = \Theta(1) \cdot \mu \cdot \text{imb}(p_i, 1 - p_i).$$

Hence $\sum_{i=1}^d \hat{f}(\{i\})^2 = \Theta(1) \cdot \mu^2 \cdot \text{tsi}(C)$. Combining this with Talagrand's lemma and rearranging yields the theorem. \square

Combining this Theorem with Proposition 4.21 implies that a typical cube drawn from $\mathcal{D}_{\text{cube}}$ has few \mathcal{V} points in it. Let us draw this conclusion for sections as well, and then strengthen it:

Proposition 4.24 *If S is a random section drawn from $\mathcal{D}_{\text{section}}$ then with probability at least $1 - \tau^2/8$, the \mathcal{V} -density of S is at most $O(1/\tau^4) \cdot \exp(-\Omega(\tau^2 d))$.*

Proof: By Proposition 4.21, $\mathbf{E}[\text{tsi}(C)] \geq \Omega(\tau^2 d)$ when C is drawn from $\mathcal{D}_{\text{cube}}$; since tsi is always at most d we conclude that a random C has $\text{tsi}(C) \geq \Omega(\tau^2 d)$ with probability at least $\Omega(\tau^2)$. In this case, Theorem 4.23 implies that C has \mathcal{V} -density at most $\exp(-\Omega(\tau^2 d))$. So we have that with probability $\Omega(\tau^2)$ a random C from $\mathcal{D}_{\text{cube}}$ has \mathcal{V} -density at most $\exp(-\Omega(\tau^2 d))$.

This implies that a random section from $\mathcal{D}_{\text{section}}$ has \mathcal{V} -density at most $2 \exp(-\Omega(\tau^2 d))$ with probability $\Omega(\tau^2)$. But in $\mathcal{D}_{\text{section}}$, sections are chosen with probability proportional to their \mathcal{V} -density (because every point is contained in the same number of sections). Thus, by a simple probabilistic argument, it follows that the average \mathcal{V} -density inside a uniformly chosen section, which is equal to the overall density of \mathcal{V} in $(\mathbb{Z}_m^d)_\infty$, is at most $(1/\Omega(\tau^2)) \cdot 2 \exp(-\Omega(\tau^2 d))$. The desired conclusion now follow from Markov's inequality. \square

We now know that a random section S almost surely has low \mathcal{V} -density. But by Proposition 4.17 a random section S has a decent chance of having most of the \mathcal{V} points of its entire tube. This gains us a factor of m :

Proposition 4.25 *If T is a random tube drawn from $\mathcal{D}_{\text{tube}}$ then with probability at least $\tau^2/8$, the \mathcal{V} -density of T is at most $(1/m) \cdot O(1/\tau^5) \cdot \exp(-\Omega(\tau^2 d))$.*

Proof: Proposition 4.17 tells us that if x is chosen randomly from \mathcal{V} and T is a random tube through x , then the expected value of $\rho(T)$ is at least τ ; hence by Markov's inequality it is at least $\tau/2$ with probability at least $\tau/2$. Furthermore, *conditioned* on $\rho(T) \geq \tau/2$, it's clear that x is in T 's section of maximum relative density with probability at least that maximum. Thus we conclude that with probability at least $\tau^2/4$ over the choice of (x, T) we have both of the following: (i) x is in T 's section of maximum relative density; and, (ii) this section accounts for at least a $\tau/2$ fraction of *all* of T 's absolute \mathcal{V} -density.

But Proposition 4.24 tells us the probability that x 's section has \mathcal{V} -density more than $O(1/\tau^4) \cdot \exp(-\Omega(\tau^2 d))$ is at most $\tau^2/8$. Hence with probability at least $\tau^2/4 - \tau^2/8 = \tau^2/8$ over the choice of (x, T) we have that T 's section of maximum relative density has \mathcal{V} -density $O(1/\tau^4) \cdot \exp(-\Omega(\tau^2 d))$ and yet accounts for at least a $\tau/2$ fraction of its overall \mathcal{V} -density. Since T has m sections, the proposition follows. \square

We may now complete the proof of Lemma 4.12. Using Proposition 4.25 and the same argument used in the second part of Proposition 4.24's proof, we conclude that the expected \mathcal{V} -density of a random tube drawn from $\mathcal{D}_{\text{tube}}$ is at most $(1/m) \cdot O(1/\tau^7) \cdot \exp(-\Omega(\tau^2 d))$. This certainly implies that if a tube is chosen *uniformly at random* from among all dm^{d-1} of $(\mathbb{Z}_m^d)_\infty$'s tubes, its expected \mathcal{V} -density is at most $(1/m) \cdot O(1/\tau^7) \cdot \exp(-\Omega(\tau^2 d))$. But this completes the proof of Lemma 4.12 because choosing a uniformly random tube and then a uniformly random point from it yields a uniformly random point in $(\mathbb{Z}_m^d)_\infty$.

References

- [ALM⁺98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, 1998.
- [AS98] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of NP. *Journal of the ACM*, 45(1):70–122, 1998.
- [BCH⁺02] Jonathan Barrett, Daniel Collins, Lucien Hardy, Adrian Kent, and Sandu Popescu. Quantum nonlocality, Bell inequalities and the memory loophole. *Physical Review A*, 66(042111), 2002.
- [Bez06] Károly Bezdek. Sphere packings revisited. *Eur. J. Combinatorics*, 27(6):864–883, 2006.
- [BKLO06] Béla Bollobás, Guy Kindler, Imre Leader, and Ryan O’Donnell. Eliminating cycles in the discrete torus. In *Proc. 7th LATIN*, pages 202–210, 2006.
- [BOGKW88] Michael Ben-Or, Shafi Goldwasser, Joe Kilian, and Avi Wigderson. Multi-prover interactive proofs: How to remove intractability. In *Proc. 20th ACM Symp. on the Theory of Computing*, pages 113–131, 1988.
- [Cho89] Jaigyoung Choe. On the existence and regularity of fundamental domains with least boundary area. *J. Differential Geom.*, 29:623–663, 1989.
- [CKK⁺05] Shuchi Chawla, Robert Krauthgamer, Ravi Kumar, Yuval Rabani, and D. Sivakumar. On the hardness of approximating Multicut and Sparsest-Cut. In *Proc. 20th IEEE Confence on Computational Complexity*, pages 144–153, 2005.
- [CS98] John H. Conway and Neil Sloane. *Sphere Packing, Lattices and Groups*. Springer-Verlag, New York, 3rd edition, 1998.
- [CSUU06] Richard Cleve, William Slofstra, Falk Unger, and Sarvagya Upadhyay. Strong parallel repetition theorem for quantum XOR proof systems, 2006. arXiv:quant-ph/0608146.
- [DMR06] Irit Dinur, Elchanan Mossel, and Oded Regev. Conditional hardness for approximate coloring. In *Proc. 38th ACM Symp. on the Theory of Computing*, pages 344–353, 2006.
- [Fei91] Uriel Feige. On the success probability of the two provers in one round proof systems. In *Proc. 6th IEEE Symp. on Structure in Complexity Theory (CCC)*, pages 116–123, 1991.
- [Fei95] Uriel Feige. Error reduction by parallel repetition — the state of the art. Available on the author’s home page, 1995.
- [FG95] Uriel Feige and Michel Goemans. Approximating the value of two prover proof systems, with applications to MAX-2SAT and MAX-DICUT. In *Proc. 3rd Israel Symp. on Theory of Comp. and Sys.*, pages 182–189, 1995.
- [FK00] Uriel Feige and Joe Kilian. Two-prover protocols — low error at affordable rates. *SIAM Journal on Computing*, 30(1):324–346, 2000. (Conference version in STOC ’94.).
- [FL92] Uriel Feige and László Lovász. Two-prover one-round proof systems, their power and their problems. In *Proc. 24th ACM Symp. on the Theory of Computing*, pages 733–744, 1992.
- [For89] Lance Fortnow. Complexity-theoretic aspects of interactive proof systems. Technical Report (Ph.D. thesis) MIT-LCS//MIT/LCS/TR-447, 1989.
- [FV02] Uriel Feige and Oleg Verbitsky. Error reduction by parallel repetition — A negative result. *Combinatorica*, 22(4):461–478, 2002.
- [GW95] Michel Goemans and David Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42:1115–1145, 1995.

- [Hal00] Thomas Hales. Cannonballs and honeycombs. *Notices of the AMS*, 47:440–449, 200.
- [Hal01] Thomas Hales. The honeycomb conjecture. *Disc. and Comp. Geom.*, 25:1–22, 2001.
- [Hås01] Johan Håstad. Some optimal inapproximability results. *Journal of the ACM*, 48(4):798–859, 2001.
- [Hol06] Thomas Holenstein. Parallel repetition: simplifications and the no-signaling case, 2006. <http://arxiv.org/abs/cs.CC/0607139>.
- [Kho02] Subhash Khot. On the power of unique 2-prover 1-round games. In *Proc. 34th ACM Symp. on the Theory of Computing*, pages 767–775, 2002.
- [KKMO04] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O’Donnell. Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? In *Proc. 45th IEEE Symp. on Foundations of Comp. Sci.*, pages 146–154, 2004.
- [KO06] Subhash Khot and Ryan O’Donnell. SDP gaps and UGC-hardness for MaxCutGain. In *Proc. 47th IEEE Symp. on Foundations of Comp. Sci.*, 2006.
- [KR03] Subhash Khot and Oded Regev. Vertex Cover might be hard to approximate to within $2 - \epsilon$. In *Proc. 18th IEEE Conference on Computational Complexity*, pages 379–386, 2003.
- [KV05] Subhash Khot and Nisheeth Vishnoi. The Unique Games Conjecture, integrality gap for cut problems and embeddability of negative type metrics into ℓ_1 . In *Proc. 46th IEEE Symp. on Foundations of Comp. Sci.*, pages 53–62, 2005.
- [Lhu81] Simon Lhuillier. Mémoire sur le minimum de cire des alvéoles des abeilles & en particulier sur un minimum minimorum relatif à cette matiere. *Nouveaux Mémoires de l’Académie Royale des Sciences et Belles-Lettres*, pages 277–300, 1781.
- [MOO05] Elchanan Mossel, Ryan O’Donnell, and Krzysztof Oleszkiewicz. Noise stability of functions with low influences: invariance and optimality. In *Proc. 46th IEEE Symp. on Foundations of Comp. Sci.*, pages 21–30, 2005.
- [MOR⁺06] Elchanan Mossel, Ryan O’Donnell, Oded Regev, Jeffrey Steif, and Benjamin Sudakov. Non-interactive correlation distillation, inhomogeneous Markov chains, and the reverse Bonami-Beckner inequality. *Israel J. of Math.*, 154:299–336, 2006.
- [Pla73] Joseph Plateau. *Statique expérimentale et théorique des liquides soumis aux seules forces moléculaires*. Gauthier-Villars, Paris, 1873.
- [Raz98] Ran Raz. A parallel repetition theorem. *SIAM Journal on Computing*, 27(3):763–803, 1998. (Conference version in STOC ’95.).
- [SM96] John Sullivan and Frank Morgan eds. Open problems in soap bubble geometry. *Intl. J. Math.*, 7:833–842, 1996.
- [SSZ04] Michael Saks, Alex Samorodnitsky, and Leonid Zosin. A lower bound on the integrality gap for Minimum Multicut in directed networks. *Combinatorica*, 24(3):525–530, 2004.
- [ST06] Alex Samorodnitsky and Luca Trevisan. Gowers uniformity, influence of variables, and PCPs. In *Proc. 38th ACM Symp. on the Theory of Computing*, pages 11–20, 2006.
- [Tal96] Michel Talagrand. How much are increasing sets positively correlated? *Combinatorica*, 16(2):243–258, 1996.
- [Tho87] William Thomson. On the division of space with minimum partitional area. *Philos. Mag. Lett.*, 24:503–514, 1887.

- [TSSW00] Luca Trevisan, Gregory Sorkin, Madhu Sudan, and David Williamson. Gadgets, approximation, and linear programming. *SIAM Journal on Computing*, 29(6):2074–2097, 2000.
- [WKC⁺04] Denis Weaire, Norbert Kern, Simon Cox, John Sullivan, and Frank Morgan. Pressures in periodic foams. *Proc. Royal Soc. London A*, 460(2042):569–579, 2004.
- [WP94] Denis Weaire and Robert Phelan. A counter-example to Kelvin’s Conjecture on minimal surfaces. *Philos. Mag. Lett.*, 69(2):107–110, 1994.

A The permutohedron

Given that the Voronoi cell of the lattice A_2^* (i.e., the regular hexagon) solves the Choe Problem for Flat Tori in \mathbb{R}^2 , and that the Voronoi cell of the lattice A_3^* (i.e., the truncated octahedron) is very close to the conjectured solution in \mathbb{R}^3 (and is conjectured [Bez06] to be the exact solution if the cell is restricted to be convex), it is natural to wonder if the Voronoi cell of the lattice A_d^* is a good solution in d dimensions. For a definition of the lattice A_d^* , see, e.g., Conway and Sloane [CS98]; for our purposes, it suffices to know that its Voronoi cell is the *permutohedron* (see [CS98, p. 474]).

The d -dimensional permutohedron, P_d , is best realized as the convex hull of the $(d+1)!$ points in \mathbb{R}^{d+1} gotten by permuting the coordinates of the vector

$$\left(-\frac{d}{2}, -\frac{d-2}{2}, -\frac{d-4}{2}, \dots, \frac{d-2}{2}, \frac{d}{2}\right).$$

This is really a d -dimensional polytope since all the points described lie in the hyperplane $\sum_{i=1}^{d+1} x_i = 0$. As examples, we have that P_1 is a line segment of length $\sqrt{2}$ and that P_2 is the regular hexagon of side length $\sqrt{2}$. It is also convenient to define P_0 to be the point with “0-dimensional volume” 1. Conway and Sloane show that the d -dimensional volume of P_d is

$$V_d = (d+1)^{d-1/2};$$

they also show that the typical face of P_d is a prism congruent to $P_r \times P_s$ for numbers $r, s \in \mathbb{N}$ satisfying $r+s = d-1$, and that P_d has precisely $\binom{d+1}{r+1}$ faces of this type.

Using this information, we can estimate the surface area of P_d ; we are not aware of any previous literature making this calculation. Writing $D = d+1$, $R = r+1$, $S = D - R$, we get that

$$\text{vol}_{d-1}(\partial P_d) = \sum_{R=1}^{D-1} \binom{D}{R} R^{R-3/2} S^{S-3/2}. \quad (6)$$

Since the summands are symmetric in R and S we can concentrate on the case $R \leq D/2$. For “small” R , $R \leq D^{1/3}$, we will use

$$\binom{D}{R} = \frac{D^R}{R!} \cdot (1 \pm O(D^{-1/3})), \quad (7)$$

which follows from elementary considerations; for “large” R , $D^{1/3} \leq R \leq D/2$, we will use

$$\binom{D}{R} = \frac{D^{D+1/2}}{\sqrt{2\pi} R^{R+1/2} S^{S+1/2}} \cdot (1 \pm O(D^{-1/3})), \quad (8)$$

which follows from Stirling’s Formula. Substituting (7) into (6) the small- R summands are estimated as

$$\frac{D^R}{R!} R^{R-3/2} (D-R)^{D-R-3/2} = V_d \cdot \frac{R^{R-3/2}}{R!} \cdot \left(1 - \frac{R}{D}\right)^{D-R-3/2} = V_d \cdot \frac{R^{R-3/2}}{R!} e^{-R}$$

up to a factor of $1 \pm O(D^{-1/3})$. Substituting (8) into (6) the large- R summands are estimated as

$$V_d \cdot \frac{1}{\sqrt{2\pi} R^2} \cdot \frac{D^2}{(D-R)^2} = V_d \cdot \frac{1}{\sqrt{2\pi} R^2}$$

up to a factor of $1 \pm O(D^{-1/3})$. We therefore conclude

$$\text{vol}_{d-1}(\partial P_d) = \left(2 \sum_{1 \leq R \leq d/2} T_R \right) \cdot V_d \cdot (1 \pm O(d^{-1/3})),$$

where

$$T_R = \begin{cases} \frac{R^{R-3/2}}{R!} e^{-R} & \text{if } R \leq d^{1/3}, \\ \frac{1}{\sqrt{2\pi R^2}} & \text{else.} \end{cases}$$

Note that the former quantity approaches the latter as $R \rightarrow \infty$, by Stirling's Formula. Indeed it is easy to see that the tail of the sum past $R = d^{1/3}$ is negligible compared to our error term, and in fact we could just as well sum the first definition of T_R over all $R \in \mathbb{N}$. In other words, we have shown the following:

Theorem A.1 *Let $K = \sum_{R=1}^{\infty} \frac{R^{R-3/2}}{R!} e^{-R} \approx .6186$. Then the surface area of the d -dimensional permutohedron P_d with volume V_d is equal to $2K \cdot V_d \cdot (1 \pm O(d^{-1/3}))$.*

Suppose we now scale P_d by a factor of α in each dimension so that its volume becomes 1. Then its surface area decreases by a factor of α/V^d , making it $2K \cdot \alpha \cdot (1 \pm O(d^{-1/3}))$. Since the appropriate α is $d - O(\log d)$, we conclude:

Corollary A.2 *When the d -dimensional permutohedron (which tiles \mathbb{R}^d by the lattice A_d^*) is scaled to have volume 1, its surface area is $2K \cdot d - O(d^{2/3})$, where $K \approx .6186$ is the constant from Theorem A.1.*

Since the d -dimensional unit cube has surface area $2d$, we see that the permutohedron foam is a better solution to the Choe Problem for Flat Tori, but only by a constant factor of K .