Spherical Cubes and Rounding in High Dimensions

Guy Kindler*
Weizmann Institute
guy.kindler@weizmann.ac.il

Ryan O'Donnell†
Carnegie Mellon University
odonnell@cs.cmu.edu

Anup Rao‡
Institute for Advanced Study
arao@ias.edu

Avi Wigderson
Institute for Advanced Study
avi@math.ias.edu

June 16, 2009

Abstract

What is the least surface area of a shape that tiles $\mathbb{R}^d$ under translations by $\mathbb{Z}^d$? Any such shape must have volume 1 and hence surface area at least that of the volume-1 ball, namely $\Omega(\sqrt{d})$. Our main result is a construction with surface area $O(\sqrt{d})$, matching the lower bound up to a constant factor of $2\sqrt{2\pi/e} \approx 3$. The best previous tile known was only slightly better than the cube, having surface area on the order of d.

We generalize this to give a construction that tiles $\mathbb{R}^d$ by translations of any full rank discrete lattice Λ with surface area $2\pi \left\|V^{-1}\right\|_{\mathrm{fb}}$, where V is the matrix of basis vectors of Λ, and $\|\cdot\|_{\mathrm{fb}}$ denotes the Frobenius norm. We show that our bounds are optimal within constant factors for rectangular lattices. Our proof is via a random tessellation process, following recent ideas of Raz [11] in the discrete setting.

Our construction gives an almost optimal noise-resistant rounding scheme to round points in $\mathbb{R}^d$ to rectangular lattice points.

1 Introduction

The d-dimensional unit cube tiles $\mathbb{R}^d$ by $\mathbb{Z}^d$. That is, its translations by vectors from $\mathbb{Z}^d$ cover $\mathbb{R}^d$, and the interiors of translations of it by different vectors from $\mathbb{Z}^d$ are disjoint.

In this paper, we consider the problem of finding a body that tiles $\mathbb{R}^d$ by $\mathbb{Z}^d$, and has the smallest possible surface area. This kind of problem is called a *foam problem*, since foams are simply tilings of space that try to minimize surface area. The best previous construction was based on the exact solution of the problem for the case of $d = 2$ [3] (Figure 1.3), and gave surface area approximately $1.93d$, only slightly better than $2d$, the surface area of the d-dimensional cube. For three or more dimensions, even potential candidates for the optimal solution are not known. In this paper we define a distribution on bodies that tile $\mathbb{R}^d$ by $\mathbb{Z}^d$ and have expected surface area at most $4\pi\sqrt{d}$. This comes close to an obvious lower bound, the surface area of a ball of volume one, which behaves asymptotically like $\sqrt{2\pi e} \cdot \sqrt{d}$. Our construction is thus asymptotically optimal up to a factor of $2\sqrt{2\pi/e}$.

Theorem 1.1. *For all $d > 0$ there exists a body which tiles $\mathbb{R}^d$ by $\mathbb{Z}^d$, and has surface area at most $4\pi\sqrt{d}$. Moreover, this body is contained in $(-1, 1)^d$, and it contains the origin.*

The ideas for our proof originate in the study of parallel repetition of *two player games*. A connection between the parallel repetition question and foams was observed recently in Feige et al. [5], where it was shown that improving the upper bounds on the success probability of the repeated odd-cycle game would imply new

*Supported by the Koshland fellowship and by the Binational Science Foundation (BSF).
†Supported in part by NSF Grant No. CCF-0747250.
‡Supported in part by NSF Grant No. CCR-0324906.

lower-bounds on the surface area of bodies tiling $\mathbb{R}^d$ by $\mathbb{Z}^d$. Subsequently, Raz [11] gave an example showing that the known upper bounds for the repeated odd-cycle game cannot be significantly improved. While it is not known that *any strategy* for the repeated odd-cycle game can be translated into a foam with small surface area, in this paper we give a continuous version of Raz's example that does give a foam with optimal surface area.

Raz's example was based crucially on a lemma of Holenstein's [6, Lemma 8], that showed a certain sampling algorithm. While Raz uses Holenstein's lemma as a black box, the main idea in our construction is to get a continuous and concrete version of the sampling algorithm of Holenstein, and use it as part of our construction.

1.1 Noise-resistant rounding

A *rounding scheme* is a random method of mapping each of the points in $\mathbb{R}^d$ to a point in $\mathbb{Z}^d$, and we say that it is *noise resistant* if the probability that close by points are rounded to different lattice points is small. Following is a formal definition.

Definition 1.2 (rounding scheme)**.** A d-dimensional rounding scheme is a distribution over functions mapping $\mathbb{R}^d$ to $\mathbb{Z}^d$. A rounding scheme is a family containing one d-dimensional rounding scheme for each dimension d. A rounding scheme is called *proper* if for some constant c, the ℓ_∞ distance between a point and its rounding is uniformly bounded by c.

For any $\delta > 0$, the δ-noise sensitivity of a rounding scheme is the maximum over all points $x, x' \in \mathbb{R}^d$ with $\|x - x'\| \leq \delta$ of the probability that the rounding of x is different from the rounding of x'.

It turns out that along with the above mentioned foam construction, our techniques give a new rounding scheme that is much better than what was previously known[1]. We think that the problem of finding noise-resistant rounding schemes is natural and interesting, and we hope it will have applications in the future.

Theorem 1.3 (Proper rounding)**.** *There exists a proper rounding scheme of $\mathbb{R}^d$ whose δ-noise sensitivity is bounded by $O(\delta + \exp(d)\delta^2)$.*

Our rounding scheme has the additional property of being periodic – each of the functions in our distribution has a period of $\mathbb{Z}^d$, so how a vector is rounded only depends on its fractional part. We observe that any proper rounding scheme must have δ noise sensitivity at least $\Omega(\delta)$: consider an axis-parallel line segment of length slightly more than 2 (its length will be the same both in ℓ_2 and in ℓ_∞ norm). On one hand, the length of the segment ensures that its endpoints are rounded to different lattice points. On the other hand, a proper rounding scheme which has δ-noise sensitivity smaller than $\delta/2$ would round both endpoints to the same lattice point with positive probability, as can be seen by breaking the segment into pieces of length at most δ (we assume $\delta \ll 1$ for convenience) and using a union bound argument.

As far as we know, no proper rounding scheme was known where the noise sensitivity is better that $\Omega(\delta\sqrt{d})$, which can be obtained just by rounding each coordinate independently. Our rounding scheme is better in the regime where $\delta \ll 2^{-d}$. As pointed out to us by Noga Alon[2], using known techniques one can get a rounding scheme whose δ-noise sensitivity is also of order $\theta(\delta)$, and while it is not proper, it still ensures that a vector in $\mathbb{R}^d$ is within ℓ_∞ distance at most $\sqrt{d}$ from its rounding.

Another somewhat similar result appears in [2], where a random partition of $\mathbb{R}^d$ into bodies of volume at most 1 and with diameter at most $O(\sqrt{d})$ is shown, such that points of distance δ end up in the same element of the partition (giving a *clustering scheme*) with probability at most $O(\delta)$. While the partition in [2] does not give rise to a proper rounding scheme, it does share some ideas with our construction.

1.2 General Lattices

Let us discuss how to generalize our results for the case of the lattice $\mathbb{Z}^d$ to arbitrary full-rank lattice in $\mathbb{R}^d$. Given any discrete, full-rank[3] lattice Λ in $\mathbb{R}^d$, we consider bodies that tile $\mathbb{R}^d$ by Λ – such a body is called a

[1] Given a foam tiling $\mathbb{R}^d$ with period $\mathbb{Z}^d$, the construction of a corresponding rounding scheme is straightforward. However the analysis of our rounding scheme requires more than just the properties of the foam stated in Theorem 1.1.

[2] The idea is to use an efficient tiling according to a well chosen volume 1 lattice Λ to round points to points in Λ. Hall's theorem then shows that there is a matching between points of Λ and Z^d such that points of distance at most twice the diameter of the tiles are matched to each other. This gives the rounding to Z^d.

[3] Throughout, we assume lattices are always full-rank.

fundamental domain and is defined formally below. To avoid technical difficulties, we want to only consider bodies that have nice smooth boundaries.

Definition 1.4. A set in $\mathbb{R}^d$ is called a $\mathcal{C}^1$ surface if it is the image of a compact set $M \subseteq \mathbb{R}^{d-1}$ under a differentiable function whose Jacobian matrix is of full rank (namely of rank $d-1$) at each point in M. A set is called piecewise C^1 if it the union of $\mathcal{C}^1$ surfaces.

Definition 1.5 (fundamental domain)**.** A compact body $K \subseteq \mathbb{R}^d$ is called a fundamental domain of a full-rank lattice Λ, if it has a piecewise $\mathcal{C}^1$ boundary, and in addition $\cup_{v\in\Lambda}(K+v) = \mathbb{R}^d$ and the interiors of the elements in this union are disjoint.

A spine of a torus. Another related object is a *spine* of the torus $\mathbb{R}^d/\Lambda$. This is a $d-1$ dimensional surface in $\mathbb{R}^d/\Lambda$ that intersects every homotopically nontrivial cycle (a closed loop that cannot be continuously deformed to a point) in the torus. In plainer words, it is a "wall" that blocks all paths that "wrap around" the torus.

We can ask the following essentially equivalent questions:

Question 1: What is the least boundary area of a fundamental domain of the torus $\mathcal{T} = \mathbb{R}^d/\Lambda$?

Question 2: What is the least surface area of a piecewise $\mathcal{C}^1$ spine in $\mathcal{T}$?

The answer to Question 1 is at most twice the answer of Question 2, since any spine can be used to get a fundamental domain whose boundary has at most twice the surface area of the spine (we omit a formal proof).

Definition 1.6. For a lattice Λ in $\mathbb{R}^d$, define

$$\mathcal{A}(\mathbb{R}^d/\Lambda) = \limsup\{|\partial\mathcal{S}| : \mathcal{S} \text{ is a piecewise } \mathcal{C}^1 \text{ spine for } \mathbb{R}^d/\Lambda\}.$$

Let us reformulate Theorem 1.1 using the new notation, where here and throughout the paper we write $\mathbb{T}^d = \mathbb{R}^d/\mathbb{Z}^d$ for the cubic torus.

Theorem 1.7. *For all d,*
$\mathcal{A}(\mathbb{T}^d) \leq 2\pi\sqrt{d}$.

We can generalize this result to other lattices. Given any basis $v_1, \ldots, v_d$ for the lattice Λ, let V denote the matrix whose rows are the basis vectors. Let $\|V\|_{\text{fb}} \stackrel{def}{=} \sqrt{\sum_{i,j} v_{i,j}^2}$ denote the Frobenius norm of V. Then we can prove the following theorem:

Theorem 1.8. *Let Λ be a volume 1 lattice in $\mathbb{R}^d$ and let V be a matrix whose rows are a basis for Λ. Then*

$$\mathcal{A}(\mathbb{R}^d/\Lambda) \leq 2\pi \left\|V^{-1}\right\|_{\text{fb}}$$

In the case where Λ is rectangular, namely if it has an orthogonal basis, we give a matching lower bound:

Theorem 1.9 (Lower Bound for Rectangular Lattices)**.** *If V is a matrix whose rows are an orthogonal basis for the lattice Λ and $S \subset \mathbb{R}/\Lambda$ is a spine, then*

$$\mathcal{A}(\mathbb{R}^d/\Lambda) \geq \left\|V^{-1}\right\|_{\text{fb}}$$

1.3 History of the problem

Foams were studied since as early as the 19th century (see [14]), they were extensively studied since by mathematicians, and they also have a huge variety of applications in physics, chemistry, and engineering (see [12] for some examples). A detailed account of the history of foam problems is thus beyond the scope of this paper.

We will, however, discuss some known upper bounds for $\mathcal{A}(\mathbb{R}^d/\Lambda)$ and some related results. But before that, let us mention an easy lower bound for $\mathcal{A}(\mathbb{R}^d/\Lambda)$. Without loss of generality assume that Λ is a volume-1 lattice. Then any fundamental domain D for Λ has volume 1 and hence must have surface area at least that of the volume-1 ball, by the Isoperimetric Inequality. Let us write κ_d for (half of) this ball's surface area, noting that $\kappa_d = \Theta(\sqrt{d})$. More precisely, we have:

Proposition 1.10.
$$\mathcal{A}(\mathbb{R}^d/\Lambda) \geq \kappa_d \sim \sqrt{\pi e/2}\sqrt{d}$$
*for any volume-*1 *lattice* Λ.

The most natural construction of a tiling of $\mathbb{R}^d$ by Λ is just to take the Voronoi cells of the points in Λ. If these cells are to have small surface area — say, $O(\sqrt{d})$ — then they should be "somewhat spherical". This leads one to consider lattices which give rise to good sphere-packings. It is not hard to show that if a volume-1 lattice has covering radius R and packing radius r then its Voronoi cells have surface area at most $(R/r)\kappa_d$. A well-known result of Butler [1] shows the existence in d dimensions of lattices with $R/r \leq 2 + o(1)$. Hence:

Proposition 1.11. *There* exist *volume-*1 *lattices in* $\mathbb{R}^d$ *satisfying*
$$\mathcal{A}(\mathbb{R}^d/\Lambda) \leq (2 + o(1))\kappa_d \leq O(\sqrt{d}).$$

Thus there exist lattices where we have a tight bound of $\Theta(\sqrt{d})$. In general lattices, however, the Voronoi cell construction can be arbitrarily far off from the $\Omega(\sqrt{d})$ lower bound. In this paper we first show that the surface area of the Voronoi cells of a lattice Λ can actually be far from the optimal $\mathcal{A}(\mathbb{R}^d/\Lambda)$: for $\Lambda = \mathbb{Z}^d$ the Voronoi cells are cubes, which have surface area d, while we show that $\mathcal{A}(\mathbb{R}^d/\Lambda) \leq \sqrt{\pi e/2}\sqrt{d}$.

Even in two dimensions, the optimal spine of $\mathbb{T}^2 = \mathbb{R}^2/\mathbb{Z}^2$ is not one that corresponds to a Voronoi cell. As proven in [3], the spine in Figure 1.3 gives the best solution. Here the fundamental domain is an "isosceles"

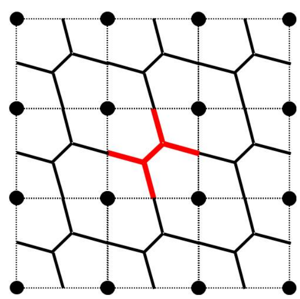

Figure 1: Optimal two dimensional tiling.

hexagon in which all angles are 120°. The spine has total length $(1 + \sqrt{3})/\sqrt{2} \approx 1.93$, slightly better that the Voronoi cell, namely the square, which gives a spine of length 2.

The question of determining the asymptotics of $\mathcal{A}(\mathbb{R}^d/\Lambda)$ was posed in Feige et al. [5], wherein special emphasis was given to the simple case of the cubic torus $\mathbb{T}^d$. Feige et al.'s interest in the problem came from showing that a "discretized" version of it plays an important role in the study of "Parallel Repetition" [4] in Complexity Theory. Feige at al. observed that by constructing prisms based on the optimal solution in $\mathbb{T}^2$ one can show
$$\mathcal{A}(\mathbb{T}^d) \leq \left(\tfrac{1+\sqrt{3}}{2\sqrt{2}} + o(1)\right) d \approx .966d,$$
very slightly improving on the trivial upper bound of d. They left as an open problem the determination the correct rate of growth, $\sqrt{d}$ vs. d. Raz [11] recently showed that $\Theta(\sqrt{d})$ is the correct rate of growth for the "discretized" version of the problem; the present paper is an extension of his result to the natural continuous case.

Although we are content to study the asymptotics of $\mathcal{A}(\mathbb{R}^d/\Lambda)$, the question of determining it precisely has also been pursued. In 1989, responding to questions of Michael Freedman, Choe [3] considered Question 1 for the case of general compact 3-manifolds. His main result was to show that there exists a fundamental domain whose surface area is minimal among those with Lipschitz boundary. He also proved optimality of the above-illustrated solution for $\mathbb{T}^2$, and gave the case of $\mathbb{T}^3$ as an open problem. As far as we know, no one has even conjectured an optimal solution for $\mathcal{A}(\mathbb{T}^3)$. In our work, we resolve this problem up to a constant factor for every d.

1.4 Subsequent work

Upon hearing a lecture on the results of this paper, Deligne asked the following natural question: "What is the minimum ratio of surface area to volume, of a body contained in $(-1,1)^d$?" As is easy to see, the analysis of one step (called the "pre-bubble") of our probabilistic tiling-construction implies the existence of a body in $(-1,1)^d$ for which this ratio is $O(\sqrt{d})$ – this ratio is optimal up to the implied constant.

Following Deligne's question, Alon and Klartag [7] have expressed this isoperimetric problem as a Dirichlet boundary problem, and showed that Cheeger's inequality and known spectral estimates directly imply the existence of such a body as well. Indeed, the appearance of the function $\Pi_i sin^2 \pi x_i$ in our sampling procedure and its optimality gets perhaps a more straightforward explanation from their viewpoint.

We note that they also proceed to give a probabilistic construction of a periodic tiling via random shifts of this body, in a similar fashion to our paper, and with a somewhat simpler analysis. Also, combining their approach with known relations between the vertex expansion of a graph and its spectral properties, they also gave similar results for some discrete graphs.

2 Proof Overview

In this section we shall reserve most of the discussion for the proof of Theorem 1.7. Suppose $A \subsetneq \mathbb{T}^d$ is an open set. One way to *define* the surface area of A is to let f be its 0-1 characteristic function and consider

$$\int \|\nabla f\|,$$

where ∇f denotes the gradient of f, and all integrals are taken over $\mathbb{T}^d$ unless otherwise specified. Of course, this does not precisely make sense, since f is not differentiable. More formally, one can take the total variation of f, or consider $\int \|\nabla f_i\|$ for sequences of smooth functions (f_i) approaching f pointwise. We can thus think of the problem of finding an open fundamental domain for $\mathbb{T}^d$ with small surface area as follows:

Task 1: Find $f : \mathbb{T}^d \to \{0,1\}$ such that:

1. $\int f = 1$.
2. The level set $\{x : f(x) = 0\}$ is a spine for $\mathbb{T}^d$.
3. $\int \|\nabla f\|$ is as small as possible.

2.1 A randomized relaxation

The first idea in the proof of Theorem 1.7 is that we may *relax* condition 1 above by taking f to be a continuous density function rather than a $0,1$-valued function (and keeping the other conditions intact). Indeed we show that given such a relaxed solution, there is a *randomized* construction of a spine with expected surface area $\int \|\nabla f\|$. Our construction will work by partitioning the $\mathbb{R}/\Lambda$ into color regions, with the guarantee that no color region contains a homotopically nontrivial cycle. Once we have such a partition, we shall argue that the union of the boundaries of the color regions form a spine. Assuming f is continuous and $M = \|f\|_\infty$, the construction is as follows:

Construction:

1. Let all points in $\mathbb{T}^d$ be "uncolored".
2. For $i = 1, 2, 3, \ldots$, until all points are colored:
3. Choose a uniformly random "translate" $Z \in \mathbb{T}^d$.
4. Choose a uniformly random "height" $T \in (0, M)$.

5. Let B_i be the "pre-bubble" $B_i = \{x \in \mathbb{T}^d : f(x - Z) > T\}$.

6. Color all uncolored points in B_i with color i. The colored points form a *bubble*.

7. Output the union of the boundaries of the color regions.

It is easy to check that with probability 1 the construction halts in finitely many rounds.

Proposition 2.1. *Assuming f is continuous, the construction halts after finitely many rounds with probability* 1.

Proof: Since f is nonnegative and $\int f = 1$, there must exist some positive $t_0 > 0$ such that $P := \{x \in \mathbb{T}^d : f(x) > t_0\}$ has positive measure $\eta > 0$. Since f is continuous, P is open, and so P contains a closed cube C of positive measure η'.

Partition $\mathbb{T}^d$ into subcubes of side length less than half that of C. We now have that there is some strictly positive $\epsilon > 0$ such that each subcube c has probability at least ϵ of being completely colored in any round of the construction. This is because there is a vol(c) chance that the random translate Z will be in c, and an independent t_0/M chance that the random height T is smaller than t_0; when both of these happen, c is completely contained in the pre-bubble defined by Z and T.

We now have a finite number of events (each subcube being completely colored in a single round), each of which occurs with some strictly positive probability in each round. It follows that all events eventually occur after finitely many rounds, with probability 1. □

The idea behind this construction comes from Raz's work [11] on the discretized version of $\mathcal{A}(\mathbb{T}^d)$; more specifically, it comes from the proof of Holenstein's Lemma [6, Lemma 8] (see also [10, Lemma 4.1]). Our analysis of it does not follow from either work, however. The construction is strongly reminiscent of random tessellation and crystallization processes; see, e.g., [9]. Also, as pointed to us by James Lee, a very similar construction appeared in [2], except that balls were used there instead of our pre-bubbles, and $\mathbb{R}^d$ was partitioned instead of $\mathbb{T}^d$.

We observe here the correctness of the construction:

Proposition 2.2. *The surface output by the construction is a spine for $\mathbb{T}^d$.*

Proof: Suppose otherwise; then there is a homotopically nontrivial loop L entirely within one bubble, i.e., color region. This L is contained in some single pre-bubble, $\{x \in \mathbb{T}^d : f(x - z) > t\}$, where $t > 0$. Hence L can be translated to a hotomopically nontrivial loop L' contained in the set $\{x \in \mathbb{T}^d : f(x) > t\}$. But f's 0-set is a spine, by assumption, and thus must intersect L'. This is a contradiction. □

In Section 3 we analyze the expected surface area of the spine produced by the construction. Let ∇f denote the vector of partial derivatives of f. Then we shall prove:

Definition 2.3. We say the function $f : \mathbb{T}^d \to \mathbb{R}^{\geq 0}$ is "nice" if it is $\mathcal{C}^2$ and has the property that ∇f has only finitely many zeros on the set $\{x : f(x) \neq 0\}$.

Theorem 2.4. *Let $f : \mathbb{T}^d \to \mathbb{R}^{\geq 0}$ satisfy $\int f = 1$. Further, assume f is "nice" (see below). Then for the above construction, the expected surface area of the boundary between bubbles is*

$$\int \|\nabla f\|.$$

Given that our construction is randomized, it is an interesting open question to come up with an explicit deterministic construction that matches its performance.

Any spine given by our construction leads to a rounding scheme in the natural way: use the spine to get a tiling of $\mathbb{R}^d$, and then round points in every body to the unique lattice point that lies in the body. The fact that the scheme obtained by our construction is proper follows from the fact that the body constructed by our scheme has ℓ_∞ diameter at most 2.

Theorem 2.5. *Let $f : \mathbb{T}^d \to \mathbb{R}^{\geq 0}$ satisfy $\int f = 1$. Further, assume f is "nice". Then for the above construction, the δ-noise sensitivity of the corresponding rounding scheme is at most*

$$\max_{u \in S^{d-1}} O\left(\delta \cdot \int |\langle \nabla f, u \rangle| + W\delta^2\right)$$

where W is an upperbound on the second derivatives of f.

2.2 Finding a good f

Given Theorem 2.4 and Theorem 1.3, we may equally well consider the more general task of finding a good "density function" f. The second idea in the proof is that we may obtain a good solution even by fixing f's 0-set to be the naive spine $\{x \in [0,1)^d : x_i = 0 \text{ for some } i\}$. Indeed, we will show that solving the following problem gives a very good solution for Theorem 2.4.

Task 2: Find a ("nice") $f : \mathbb{T}^d \to \mathbb{R}^{\geq 0}$ such that:

1. $\int f = 1$.
2. $f(x) = 0$ if $x_i = 0$ for some i.
3. $\int \|\nabla f\|$ is as small as possible.

In Task 2, the presence of $\|\nabla f\|$ is analytically difficult. We can make it more tractable by expressing $f = g^2$. Then we have the constraint $\int g^2 = 1$, and

$$\begin{aligned}\int \|\nabla f\| = 2\int |g| \cdot \|\nabla g\| &\leq 2\sqrt{\int g^2}\sqrt{\int \|\nabla g\|^2} \\ &= 2\sqrt{\int \|\nabla g\|^2},\end{aligned} \tag{1}$$

where we used Cauchy-Schwarz. This helps because $\int \|\nabla g\|^2$ is easier to work with. It remains to analyze the following:

Task 3: Find $g : \mathbb{T}^d \to \mathbb{R}$ such that:

1. $\int g^2 = 1$.
2. $g(x) = 0$ if $x_i = 0$ for some i.
3. $2\sqrt{\int \|\nabla g\|^2}$ is as small as possible.

So far these proof ideas all have analogues in Raz's work. We give an improvement by solving Task 3 optimally:

Theorem 2.6. *The minimizing g for Task 3, among piecewise $\mathcal{C}^1$ functions, is*

$$g(x) = \prod_{i=1}^{d} \sqrt{2}\sin(\pi x_i).$$

For this function, $2\sqrt{\int \|\nabla g\|^2} = 2\pi\sqrt{d}$ (and also $f = g^2$ is "nice").

The proof is an elementary use of the Fourier series and is given in Section 4. We note that the maximum value obtained by the induced density function f is $O(\exp(d))$. The expected volume of a pre-bubble chosen according to this density function is $\exp(-d)$.

With regards to finding a proper rounding, it turns out that we can again use the Cauchy-Schwartz inequality to get an upperbound:

Theorem 2.7. *Set*

$$f(x) = \prod_{i=1}^{d} 2\sin^2(\pi x_i).$$

Then, $\max_{u \in S^{d-1}} \int |\langle \nabla f, u\rangle| = O(1)$*, and the second derivatives of* f *are* $O(\exp(d))$*.*

2.3 Completing the proof

The proof of Theorem 1.7 follows immediately from Theorems 2.4 and 2.6 and the Cauchy-Schwarz argument (1). An illustration of the construction in $\mathbb{T}^2$ with $f(x,y) = \sin^2(\pi x)\sin^2(\pi y)$ appears below.

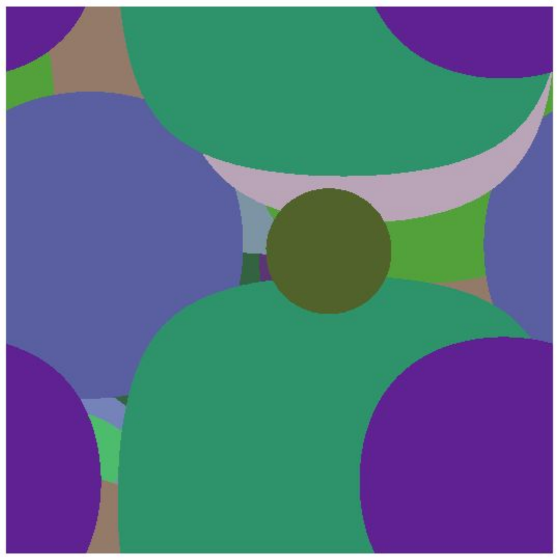

Similarly, the proof of Theorem 1.3 follows from Theorem 2.5 and Theorem 2.7.

The proof of Theorem 1.8 for $\mathbb{R}^d/V\mathbb{Z}^d$ follows by applying the linear transformation V to the random spine constructed for $\mathbb{T}^d$. The analysis is given in Section 5.

3 Analyzing the Construction in terms of f

In this section we prove Theorem 2.4. Assume throughout this section that f is "nice" as in Definition 2.3. We begin with a straightforward observation:

Proposition 3.1. *The spine output by the construction is a* $(d-1)$*-dimensional surface which is piecewise* $\mathcal{C}^1$*.*

Proof: It suffices to show this is true of the boundary of each pre-bubble. Since f is continuous, each pre-bubble has boundary $B = f^{-1}(t)$ for some $0 < t < M$. By "niceness", $\nabla f = 0$ on at most finitely many points of B. Finally, since f is $\mathcal{C}^1$, a general version of the Implicit Function Theorem implies that B is a piecewise $\mathcal{C}^1$ $(d-1)$-dimensional surface. □

Given a piecewise $\mathcal{C}^1$ surface in $\mathbb{T}^d$, we can express its area via a "Buffon's Needle" or Cauchy-Crofton-type formula.

Definition 3.2. For a point $a \in \mathbb{T}^d$ and a direction $u \in S^{d-1}$, we define the "needle" (line segment) $\ell_{a,\delta u}$ of length $0 < \delta < 1$ to be $\{a + \lambda u : \lambda \in [0, \delta]\} \subset \mathbb{T}^d$.

The following result is from the Integral Geometry textbook of Santaló [13] (the $d = 2$ case is stated as (8.11) therein; the extension to d dimensions is discussed on page 274):

Theorem 3.3. *Let S be a piecewise $\mathcal{C}^\infty$ surface in $\mathbb{T}^d$. Let $\ell_{a,\delta u}$ be a uniformly random needle of length δ; i.e., $a \in \mathbb{T}^d$ and $u \in S^{d-1}$ are chosen uniformly and independently. Then*

$$\mathop{\mathbf{E}}_{a,u}\big[\#(\ell_{a,\delta u} \cap S)\big] = C_d \cdot \delta \cdot \mathrm{area}(S),$$

where $C_d \approx 1/\sqrt{d}$ is the dimension-dependent constant

$$C_d = \mathop{\mathbf{E}}_{v \in S^{d-1}}[|v_1|]$$

and $\#(\ell_{a,\delta u} \cap S)$ denotes the number of points of intersection between the needle and the surface.

Our plan is to fix a short needle $\ell_{a,\delta u}$ and estimate the expected number of intersections it makes with the random spine. The main technical theorem we need is:

Theorem 3.4. *Fix a needle $\ell = \ell_{a,\delta u}$ of length δ in $\mathbb{T}^d$. Let N be the random variable denoting the number of intersections ℓ makes with the spine S output by the construction, $\#(\ell \cap S)$. Let $W < \infty$ be an upper bound on the magnitude of f's second-order partial derivatives (recall that f is $\mathcal{C}^2$). Then, if $\beta = \int |\langle \nabla f, u\rangle| + W\delta$ is such that $\beta < 1/\delta$,*

$$\mathop{\mathbf{E}}_{S}[N] \le \frac{\delta\beta}{1 - \delta\beta}.$$

As should be expected, $\mathbf{E}[N]$ does not depend on a: the construction is translation-invariant in $\mathbb{T}^d$. Given Theorem 3.4, our main theorems Theorem 2.4 and Theorem 2.5 follow easily:

Proof: (Theorem 2.4) Let S denote the random spine output by the construction. Let β be as in Theorem 3.4. Then for δ small enough, by Theorem 3.3,

$$\begin{aligned} C_d \cdot \delta \cdot \mathbf{E}[\mathrm{area}(S)] &= \mathop{\mathbf{E}}_{S}\left[\mathop{\mathbf{E}}_{a,u}\big[\#(\ell_{a,\delta u} \cap S)\big]\right] \\ &= \mathop{\mathbf{E}}_{a,u}\left[\mathop{\mathbf{E}}_{S}\big[\#(\ell_{a,\delta u} \cap S)\big]\right] \\ &\le \mathop{\mathbf{E}}_{a,u}\left[\frac{\delta\beta}{1-\delta\beta}\right] \end{aligned}$$

Taking $\delta \to 0$, this gives us:

$$\begin{aligned} \mathbf{E}[\mathrm{area}(S)] &\le (1/C_d) \lim_{\delta \to 0} \mathop{\mathbf{E}}_{a,u} \beta \\ &= (1/C_d) \mathop{\mathbf{E}}_{u}\left[\int |\langle \nabla f, u\rangle|\right] \\ &= (1/C_d) \int \mathop{\mathbf{E}}_{u}\big[|\langle \nabla f, u\rangle|\big] \\ &= \int \|\nabla f\| \end{aligned}$$

□

Proof: (Theorem 2.5) Let $\ell_{a,\delta u}$ be a needle of length δ. The probability that the end points of this needle are rounded to different points by the rounding scheme is bounded by $\mathbf{E}[N]$, which, by Theorem 3.4, is at most $\frac{\delta\beta}{1-\delta\beta} = O(\delta\beta) = O(\delta \int |\langle \nabla f, u\rangle| + W\delta^2)$. This completes the proof of the theorem. □

Thus it remains to prove Theorem 3.4. The theorem follows immediately from Lemma 3.5 and Lemma 3.6 below.

Recall that the construction defines pre-bubbles $B_1, B_2, \ldots$. Let E_i denote the event $B_i \cap \ell \neq \emptyset$ and let M_i denote the random variable $\#(\ell \cap \partial B_i)$.

Lemma 3.5. *Let $\kappa = \mathbf{E}[M_1 \mid E_1]$. Then*

$$\kappa \leq \delta \cdot \int |\langle \nabla f, u \rangle| + W\delta^2$$

Lemma 3.6. *Let κ be as in Lemma 3.6, and assume $\kappa < 1$. Then*

$$\kappa \leq \mathop{\mathbf{E}}_{S}[N] \leq \kappa/(1-\kappa).$$

Proof: (Lemma 3.5) For completeness we begin by noting that $\mathbf{Pr}[E_1]$ is easily seen to be positive, as follows from Proposition 2.1.
Recall that we have fixed a needle $\ell = \ell_{a,\delta u}$ of length δ in $\mathbb{T}^d$. Given $z \in \mathbb{T}^d$ we will let $f_z : [0, \delta] \to \mathbb{R}^{\geq 0}$ denote the restriction of the function $f(x - z)$ to the needle ℓ. By definition,

$$E_1 = \{(z, t) : t < \|f_z\|_\infty\}$$
$$M_1 = \#\{\lambda \in [0, \delta] : f_{Z_1}(\lambda) = T_1\},$$

and hence

$$\begin{aligned}\kappa &= \mathbf{E}[M_1 \mid E_1] \\ &= \frac{\int_{\mathbb{T}^d} \int_0^{\|f_z\|_\infty} \#\{\lambda \in [0, \delta] : f_z(\lambda) = t\}\, dt\, dz}{\int_{\mathbb{T}^d} \|f_z\|_\infty\, dz}. \end{aligned} \tag{2}$$

Let's estimate the quantities in (2). First, we have

$$\int_{\mathbb{T}^d} \|f_z\|_\infty\, dz \geq \int_{\mathbb{T}^d} f(a - z)\, dz = 1 \tag{3}$$

As for the main integrand in (2), using the fact that f_z is $\mathcal{C}^1$ we have

$$\int_0^{\|f_z\|_\infty} \#\{\lambda \in [0, \delta] : f_z(\lambda) = t\}\, dt = \int_0^\delta |f_z'(\lambda)|\, d\lambda.$$

This follows easily from considering the contribution of f_z from small segments.
Since W bounds f_z's order-2 partial derivatives, we conclude that on the range $[0, \delta]$,

$$|f_z'(\lambda)| \leq |\langle \nabla f_z(a), u \rangle| + W\delta.$$

Thus we have

$$\begin{aligned}&\int_{\mathbb{T}^d} \int_0^{\|f_z\|_\infty} \#\{\lambda \in [0, \delta] : f_z(\lambda) = t\}\, dt\, dz \\ &\leq \int_{\mathbb{T}^d} \delta \cdot |\langle \nabla f_z(a), u \rangle|\, dz + W\delta^2. \end{aligned} \tag{4}$$

Combining (3) and (4) we conclude

$$\kappa \leq \delta \int_{\mathbb{T}^d} |\langle \nabla f_z(a), u \rangle|\, dz + W\delta^2.$$

Since the above integral does not depend on a, we get the claim of the lemma. □

Proof: (Lemma 3.6) Recall that M_i denotes the random variable $\#(\ell \cap \partial B_j)$. Let C_i be the event that B_i completely encloses the needle, so $C_i = E_i \wedge (M_i = 0)$. If $\cup C_i$ has not occurred after the construction ends, continue choosing pre-bubbles until it does. Since $\mathbf{Pr}[E_i] > 0$ and $\mathbf{E}[N_i \mid E_i] = \kappa < 1$, each event C_i has positive probability and therefore $\cup C_i$ will occur after finitely many pre-bubbles, with probability 1. Let R denote the least index such that C_R occurs.

Let $M'_{j_1}, M'_{j_2}, \ldots, M'_{j_K}$ denote the values of M_i for those i such that E_i occurs, up until $M'_{j_K} = 0$, i.e., $j_K = R$. We claim

$$M'_{j_1} \leq N \leq \sum_{k=1}^{K} M'_{j_k}. \tag{5}$$

The lower bound simply says that the needle has at least as many spine-intersections as it has intersections with the first pre-bubble that touches it. The upper bound holds because once a pre-bubble completely encloses the needle it will never make any more intersections with the spine, and because counting $\sum \#(\ell \cap \partial B_i)$ can only overcount $\#(\ell \cap S)$.

The distribution of each M'_{j_k} is that of $M_1 \mid E_1$ and hence $\mathbf{E}[M'_{j_k}] = \kappa$. Thus if we take expectations in (5) we get

$$\kappa \leq \mathbf{E}[N] \leq \mathbf{E}[K]\kappa,$$

using Wald's Theorem in the upper bound. Now K is distributed as the least index for which a sequence of i.i.d. random variables, $M'_{j_1}, \ldots, M'_{j_K}$, is 0. Since M'_{j_k} is integer-valued, the probability it is 0 is at least $1 - \mathbf{E}[M'_{j_K}] = 1 - \kappa$. Hence $\mathbf{E}[K] \leq 1/(1-\kappa)$, the mean of a geometric random variable with parameter $1-\kappa$. The proof is complete. □

4 Finding a good density f

In this section we prove Theorem 2.6 and Theorem 2.7.
Suppose that $g : \mathbb{T}^d \to \mathbb{R}$ is piecewise $\mathcal{C}^1$, $\int g^2 = 1$, and $g(x) = 0$ whenever $x_i = 0$ for some i. We shall first show that the minimum possible value for $\int \|\nabla g\|^2$ is $\pi^2 d$, and occurs when

$$g(x) = \prod_{i=1}^{d} \sqrt{2}\sin(\pi x_i). \tag{6}$$

Having shown this we only need to check that $f = g^2$ is "nice" in the sense of Definition 2.3.

4.1 Optimizing g for Surface Area

Expand $g : [0,1)^d \to \mathbb{R}$ in terms of its (multidimensional) Fourier sine series:

$$g(x) = \sum_{\omega \in \mathbb{N}^d} \hat{g}(\omega) \prod_{i=1}^{d} \sqrt{2}\sin(\pi\omega_i x_i), \tag{7}$$

where

$$\hat{g}(\omega) = \int g(x) \prod_{i=1}^{d} \sqrt{2}\sin(\pi\omega_i x_i).$$

We remark that we have pointwise convergence everywhere in (7), since g is piecewise $\mathcal{C}^1$ and satisfies $g(x) = 0$ whenever $x_i \in \{0,1\}$ (hence g's odd extension is continuous). More crucially, these conditions also justify term-by-term differentiation of g's sine series. Let D_j denote the jth partial derivative. Then we get the expansion

$$D_j g(x) = \sum_{\omega \in \mathbb{N}^d} \pi\omega_j \hat{g}(\omega)(\sqrt{2}\cos(\pi\omega_j x_j)) \prod_{i \neq j} \sqrt{2}\sin(\pi\omega_i x_i). \tag{8}$$

We now apply Parseval's Theorem for cosine series and sine series to both (7) and (8), obtaining

$$1 = \int g^2 = \sum_{\omega \in \mathbb{N}^d} \hat{g}(\omega)^2, \quad \text{and} \tag{9}$$

$$\int \|\nabla g\|^2 = \sum_{j=1}^{d} \sum_{\omega \in \mathbb{N}^d} \pi^2 \omega_j^2 \hat{g}(\omega)^2 = \pi^2 \sum_{\omega \in \mathbb{N}^d} \|\omega\|^2 \hat{g}(\omega)^2. \tag{10}$$

It's now clear that subject to (9), the expression in (10) is minimized when the Fourier sine spectrum is concentrated on the frequency ω with minimal $\|\omega\|$, namely $\omega = (1, \ldots, 1)$. Hence (6) is indeed the minimizer, as claimed, and the minimal value is $\pi^2 d$.

4.2 Bounding Noise Sensitivity using f

To prove Theorem 2.7, for every $u \in S^{d-1}$, we need to bound

$$\int |\langle \nabla f, u\rangle| \leq \sqrt{\int \langle \nabla f, u\rangle^2},$$

where the inequality is by applying Cauchy-Schwartz.

Let, $\rho \in \{-1, +1\}^d$ be a uniformly random vector. Then we have the following derivation:

$$\begin{aligned}
\int \langle \nabla f, u\rangle^2 &= \int \left(\sum_{i=1}^{d} u_i 4 \sin(\pi x_i) \cos(\pi x_i) \prod_{j \neq i} 2 \sin^2(\pi x_j) \right)^2 \\
&= \int \mathbf{E}_\rho \left(\sum_{i=1}^{d} u_i 4 \sin(\pi \rho_i x_i) \cos(\pi \rho_i x_i) \prod_{j \neq i} 2 \sin^2(\pi \rho_j x_j) \right)^2 \\
&= \int \mathbf{E}_\rho \left(\sum_{i=1}^{d} u_i 4 \rho_i \sin(\pi x_i) \cos(\pi x_i) \prod_{j \neq i} 2 \sin^2(\pi x_j) \right)^2 \\
&= \int \sum_{i=1}^{d} u_i^2 16 \sin^2(\pi x_i) \cos^2(\pi x_i) \prod_{j \neq i} 16 \sin^4(\pi x_j)
\end{aligned}$$

The last equality follows from the fact that all the cross terms vanish under expectation. This integral is easily seen to be bounded by $O(\sum_i u_i^2) = O(1)$, which is optimal, by the lowerbound on the δ-noise sensitivity of any rounding scheme, as discussed in the introduction.

For this choice of f, the second derivatives f f are bounded by $2^d \mathrm{poly}(d) = O(\exp(d))$, which gives us the bound we proved on the noise sensitivity.

4.3 f is "nice"

Note that g itself is not even globally $\mathcal{C}^1$ as a function on the torus $\mathbb{T}^d$; it has kinks on its 0-set, since $\sin(\pi x)$ is naturally periodic on $[-1, 1]$ rather than $[0, 1]$. Nevertheless, a trigonometric identity implies

$$(\sqrt{2} \sin(\pi x))^2 = 1 - \cos(2\pi x),$$

and this is $\mathcal{C}^\infty$ on the circle $\mathbb{T}$. Hence f is $\mathcal{C}^\infty$ and hence $\mathcal{C}^2$ on $\mathbb{T}^d$.

Next, the set $\{x : f(x) \neq 0\}$, on which we need to consider the zeros of ∇f, is clearly $(0, 1)^d$. We calculate that

$$D_j f(x) = 2^d 2\pi \sin(2\pi x_j) \cdot \prod_{i \neq j} \sin(\pi x_i),$$

from which it follows that the only zero of ∇f on $(0, 1)^d$ is at $(1/2, \ldots, 1/2)$. Hence we only have finitely many zeros, as required for "niceness".

5 General lattices

In this section we consider the problem for other volume 1 lattices Λ besides $\mathbb{Z}^d$. Let $v_1, \ldots, v_d$ denote a basis for Λ, and arrange these vectors as columns in a matrix V. Let $V^* = (V^{-1})^\top$, the matrix of dual basis vectors.

A natural way to construct a spine of low surface area for $\mathbb{R}^d/\Lambda$ is simply to take our construction for $\mathbb{R}^d/\mathbb{Z}^d$ and apply the linear transformation V. It's clear that this indeed gives a spine. Regarding its surface area:

Theorem 5.1. *The expected surface area of the spine formed in $\mathbb{R}^d/\Lambda$ by running our construction and applying the linear transformation V is*

$$\int_{\mathbb{T}^d} \|V^* \nabla f\|.$$

Proof: Although we stated Santaló's Theorem 3.3 for $\mathbb{T}^d$, in fact it holds for any volume 1 lattice, so long as the needle is short enough to fit completely inside the fundamental parallelepiped. Since we take $\delta \to 0$, this is not a concern.

Getting an analogue of Theorem 3.4 is easy. Instead of fixing a needle $\ell = \ell_{a,\delta u}$ in $\mathbb{R}^d/\Lambda$, choosing S via the construction, and then looking at the expected value of $\#(\ell \cap VS)$, we can instead fix the preimage of the needle $V^{-1}\ell$ in $\mathbb{T}^d$ and look at the expected value of $\#(V^{-1}\ell \cap S)$. Theorem 3.4 tells us this quantity equals

$$\delta \cdot \int_{\mathbb{T}^d} |\langle \nabla f, V^{-1}u \rangle| = \delta \cdot \int_{\mathbb{T}^d} |\langle V^* \nabla f, u \rangle|$$

up to $O(W^2\delta^2)$. The remainder of the proof is unchanged. $\square$

We again use the Cauchy-Schwarz argument (1) to upper-bound

$$\int_{\mathbb{T}^d} \|V^* \nabla f\| \leq 2\sqrt{\int_{\mathbb{T}^d} \|V^* \nabla g\|^2}.$$

Finally, with our choice of g from (6), it is easy to see from (8) that

$$\int_{\mathbb{T}^d} \|V^* \nabla g\|^2 = \pi^2 \sum_{i,j=1}^{d} (V^*_{ij})^2 = \pi^2 \|V^*\|_{\text{fb}}^2 = \pi^2 \left\|V^{-1}\right\|_{\text{fb}}^2.$$

Thus we get a spine for $\mathbb{R}^d/\Lambda$ whose expected surface area is at most $2\pi \left\|V^{-1}\right\|_{\text{fb}}$, completing the proof of Theorem 1.8.

6 Lower Bounds

We have already observed an $\Omega(\sqrt{d})$ lower bound on the surface area of any spine of $\mathbb{R}/\mathbb{Z}^d$ via the Isoperimetric Inequality. In this section we generalize this to give a simple lower bound (Theorem 1.9) that applies to the surface area of a spine of $\mathbb{R}/\Lambda$ for any volume 1 *orthogonal* lattice Λ. A lattice is orthogonal if it has an orthogonal basis.

The theorem follows from the following simple generalization of Pythagoras's theorem:

Theorem 6.1. *Let $v_1, \ldots, v_d$ be orthogonal vectors. For each i, let F_i denote the $d-1$ dimensional facet whose corners are the origin and all basis vectors not equal to v_i. Let S be any piecewise continuous $d-1$ dimensional manifold such that for every i, the projection of S to F_i covers F_i. Then* $\text{area}(S)^2 = \sum_{i=1}^d \text{area}(F_i)^2$.

Next, we prove Theorem 1.9:

Proof: (Proof of Theorem 1.9) Since the Frobenius norm is preserved under unitary transformations, it suffices to prove the theorem for the case that matrix V is a diagonal matrix. Let the diagonal entries be $\alpha_1, \ldots, \alpha_d$. Then note that for every i, $\text{area}(F_i) = \prod_{j \neq i} \alpha_j = 1/\alpha_i$. The last inequality follows from the fact that $\det(V) = 1$. On the other hand, V^{-1} is simply the diagonal matrix with $1/\alpha_i$ on the diagonal. Thus the square of the area of the spine is at least $\sum_{i=1}^d 1/\alpha_i^2 = \left\|V^{-1}\right\|_{\text{fb}}^2$. $\square$

7 Acknowledgments

We thank N. Alon, K. Ball, J. Håstad, L. Hoffman, A. Klivans, D. Micciancio, V. Milman, J. Lee and O. Regev for useful discussion and comments.

References

[1] G. Butler. Simultaneous packing and covering in Euclidean space. *Proc. London Math. Soc.*, 3(4):721–735, 1972.

[2] M. Charikar, C. Chekuri, A. Goel, S. Guha, and S. Plotkin. Approximating a finite metric by a small number of tree metrics. In *Proceedings of the 39th Annual Symposium on Foundations of Computer Science(FOCS-98)*, pages 379–388, Los Alamitos, CA, Nov.8–11 1998. IEEE Computer Society.

[3] J. Choe. On the existence and regularity of fundamental domains with least boundary area. *J. Differential Geom.*, 29:623–663, 1989.

[4] U. Feige. Error reduction by parallel repetition — the state of the art. Available on the author's home page, 1995.

[5] U. Feige, G. Kindler, and R. O'Donnell. Understanding parallel repetition requires understanding foams. In *Proc. 22nd IEEE Conference on Computational Complexity*, pages 179–192, 2007.

[6] T. Holenstein. Parallel repetition: simplifications and the no-signaling case. In *Proc. 39th ACM Symp. on the Theory of Computing*, pages 411–419, 2007.

[7] B. Klartag and N. Alon. Economical toric spines via cheeger's inequality. To appear, 2008.

[8] R. Kusner and J. Sullivan. Comparing the Weaire-Phelan equal-volume foam to Kelvin's foam. *Forma*, 11(3):233–242, 1996.

[9] J. Møller. Random Johnson-Mehl tessellations. *Adv. Appl. Prob.*, 24(4):814–844, 1992.

[10] A. Rao. Parallel Repetition in projection games and a concentration bound. In *Proc. 40th ACM Symp. on the Theory of Computing*, 2008.

[11] R. Raz. A counterexample to strong parallel repetition. Manuscript, 2008.

[12] J. Sadoc and N. Rivier. *Foams and emulsions. Proceedings of the NATO ASI on foams and emulsions, emulsions and cellular materials held in Cargèse, France, May 12–24, 1997.* 1999.

[13] L. Santaló. *Integral Geometry and Geometric Probability.* Cambridge University Press, 2004.

[14] Lord Kelvin (Sir W. Thomson). On the division of space with minimum partitional area. *Philosophical Magazine*, 24(151):503, 1887.

[15] D. Weaire. *The Kelvin Problem: Foam Structures of Minimal Surface Area.* CRC Press, 1997.