

Spherical Cubes: Optimal Foams from Computational Hardness Amplification

By Guy Kindler, Anup Rao, Ryan O'Donnell, and Avi Wigderson

Abstract

Foam problems are about how to best partition space into bubbles of minimal surface area. We investigate the case where one unit-volume bubble is required to tile d -dimensional space in a periodic fashion according to the standard, cubical lattice. While a cube requires surface area $2d$, we construct such a bubble having surface area very close to that of a sphere; that is, proportional to \sqrt{d} (the minimum possible even without the constraint of being periodic). Our method for constructing this “spherical cube” is inspired by foundational questions in the theory of computation related to the concept of hardness amplification. Our methods give new algorithms for “coordinated discretization” of high-dimensional data points, which have near-optimal noise resistance. We also provide the most efficient known cubical foam in three dimensions.

1. INTRODUCTION

A foam in the d -dimensional space \mathbb{R}^d is a partition of \mathbb{R}^d into bounded sets called bubbles. In such a foam, the bubbles are said to *tile* the space. The main question studied in this work is if a foam in \mathbb{R}^d has only bubbles with a given volume, what is the minimal possible average surface area of its bubbles? This fundamental question has been a focus of study for scientists in many disciplines, from physicists studying soap bubbles,²¹ to chemists studying crystal structures,¹² biologists studying cell aggregation,³ mathematicians studying sphere-packings,¹⁴ materials scientists studying polymers,²⁰ and even artists and architects.⁶ In this work, we present a new approach to the construction of tiling shapes, based on methods from computer science. This approach leads to an asymptotically optimal solution of the Cubical Foam Problem, defined below.

1.1. Foams

Questions about minimal surface area tilings of space have a very long history. In the 19th century, Thomson (Lord Kelvin) introduced the Kelvin Foam Problem,²⁶ which asks how three-dimensional space can be partitioned into bubbles of volume 1 such that the average surface area of the bubbles in the foam is minimized. This question is motivated not only by its mathematical appeal, but also by interest in the physics of foams in nature, since surface tension makes bubbles seek to minimize their surface area.

One way to design foams with small surface area is to first construct a lattice of periodically arranged points, and

then to take the Voronoi cells around each lattice point. The Voronoi cell of a lattice point x is the bubble, which includes all points that are closer to x than to any other lattice point. The solution Kelvin proposed in 1887 for his problem was based on the Voronoi foam associated with the so-called *body-centered* cubic lattice. The bubbles in this foam have a surface area $(3/4)\sqrt[3]{120\sqrt{3}+148} \approx 5.315$. Kelvin further suggested letting this foam relax, so that it conforms with Plateau's Rules for soap bubbles²¹; modern computer simulations show that this decreases the surface area to about 5.306.^{7, 18} In 1994, Weaire and Phelan²⁷ exhibited a foam with an improved average surface area of about 5.288. The Weaire–Phelan foam is formed by relaxing the Voronoi foam for a periodic subset of lattice points (Figure 1). Weaire and Phelan used the crystal structure of a certain silicon–sodium clathrate to choose the points. It is still unknown whether or not their foam optimally solves Kelvin's problem.

It is natural to study the Kelvin Foam Problem in dimensions other than three. In two dimensions, it was long believed that the best solution is to tile space with regular hexagons, which is the Voronoi foam of the triangular lattice. The optimality of this foam was conjectured as far back as in the 4th century by Pappus of Alexandria, but a mathematical proof was found only in 1999, by Hales.¹³ In higher dimensions, a lower bound on the average surface area follows from the *Isoperimetric Inequality*: the surface area of any bubble of volume 1 must be at least as large as that of a ball of volume 1. As the number of dimensions d grows, this lower bound asymptotically approaches $\sqrt{2\pi ed}$. An upper bound that matches this lower bound up to a factor of 2 can be obtained by taking the Voronoi foam of a d -dimensional lattice in which the covering-radius to packing-radius ratio tends to 2. Such a lattice can be obtained by a probabilistic construction.⁸ Hence, the minimum surface area in the d -dimensional Kelvin Foam Problem grows in proportion to the square-root of the dimension.

In our work, we consider tilings that are periodic with respect to the *integer lattice* (also known as the cubic lattice).

The original version of this paper is entitled “Spherical Cubes and Rounding in High Dimensions” and was published in the *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science*, IEEE Computer Society, 2008.

This lattice consists of the points in d -dimensional space whose coordinates are all integers. We address the following question:

Cubical Foam Problem: What is the least surface area of a bubble that partitions d -dimensional space periodically according to the integer lattice?

The Voronoi foam for the integer lattice consists of cubes of side length 1. In d dimensions, these cubes have surface area $2d$. This grows linearly with the dimension, much higher than the known lower bound of \sqrt{d} . Are there more “spherical” cubes, which still tile by the integer lattice, but have surface area closer to that of a ball? This is the main question that we answer in this work.

The Cubical Foam Problem seems to have been first formally raised by Choe.¹⁰ Choe showed that in two dimensions, the unit square whose surface area (perimeter) is 4 is not the optimal solution. Rather, the optimal solution is the isosceles hexagon shown in Figure 2, with 120° angles, side lengths $\frac{1}{\sqrt{2}} - \frac{1}{\sqrt{6}}$ and $\sqrt{\frac{2}{3}}$, and perimeter about 3.864. Choe gave the three-dimensional version as an open problem. Prior to our work, the best known solution was simply to add depth to the Choe hexagon, transforming it into the prism shown in Figure 2, with surface area 5.864.¹¹

The high-dimensional version of the Cubical Foam Problem was raised by Feige et al.¹¹ in 2007, who noted a surprising connection to a certain problem in theoretical computer science about computational hardness amplification. We shall explain the details of this connection later. A subsequent result of Raz²⁴ on the limits of such amplification, using an idea from a related paper of Holenstein,¹⁶ provided us with the tools to solve the high-dimensional Cubical Foam Problem.

1.2. Our results

- We give a probabilistic construction proving the existence of a bubble that partitions d -dimensional space according to the cubic lattice and whose surface area is at most $4\pi\sqrt{d}$. Thus, our bubble is nearly spherical, in the sense that its surface area is larger than that of a sphere by only a constant multiplicative factor (about 3.04). The best previous constructions had surface area proportional to d (like the cube itself has). We conclude that the optimal solution to the Cubical Foam Problem has surface area proportional to the square-root of the dimension, just as in the more general Kelvin Foam Problem. Thus in high dimensions, integer-lattice til-

Figure 1. (Left) Four bubbles in the Kelvin Foam, formed by relaxing the Voronoi cells of the body-centered cubic lattice. (Right) Seven bubbles in the Weaire–Phelan Foam, formed by relaxing the Voronoi cells of the A15 Packing.

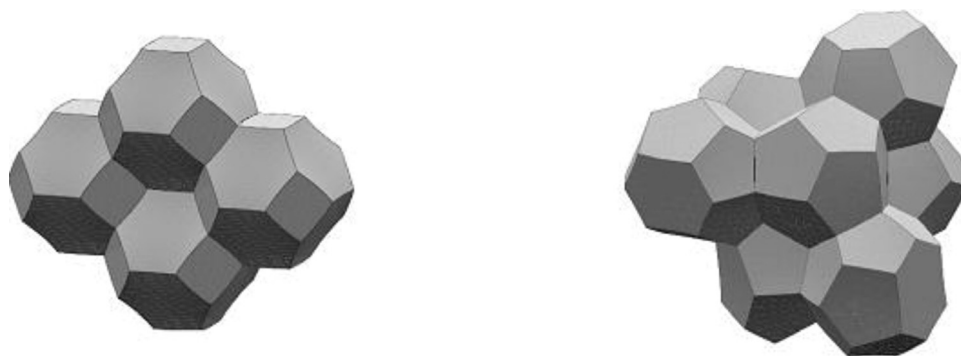



Figure 2. (Left) The Choe Hexagon, Choe’s optimal solution to the Cubical Foam Problem in two dimensions. (Right) The hexagons extruded into three-dimensional prisms. The resulting three-dimensional cubical foam is not optimal; a solution with smaller average surface area is presented in this work.



ings are essentially as efficient as arbitrary tilings.

- We show that our construction also yields a highly noise-resistant procedure for the rounding high-dimensional data. Specifically, we obtain a randomized procedure, which assigns each vector of real numbers $x = (x_1, \dots, x_d)$ to a vector of integers $n = (n_1, \dots, n_d)$, with the following two guarantees. Each n_i is always simply x_i rounded up or down to its floor or ceiling integer, yet if two real vectors x, y are at Euclidean distance ρ , then our procedure assigns them to the same integer vector except with probability at most $2 \cdot \pi \cdot \rho$. Thus, if two parties get two noisy versions x, y of the same integer vector n , they will round them to the same vector with very high probability. Somewhat remarkably, the error in this natural coordination task does not depend on the dimension at all, and depends only linearly on the noise. Previously known procedures either rounded points to far off points, so that $|x_i - n_i|$ could be as large as \sqrt{d} , or separated points at distance ρ with probability as high as $\sqrt{d}\rho$.^{9,19}
- Using different (and ad hoc) methods, we give an explicit construction for the Cubical Foam Problem in three dimensions with a surface area about 5.602. This beats the surface area of the prism obtained from Choe's prism by about 4.5%.

1.3. Computational hardness amplification

Our method for constructing an efficient cubical foam has a surprising inspiration: the subject of computational hardness amplification in the study of the theory of computation. Consider a computational task T , such as solving a system of equations, finding the best move in a chess position, optimizing a schedule under constraints, etc. The difficulty, or hardness, of T is measured in terms of the computational resources (say the number of steps in an algorithm) required to obtain a solution of a given quality. *Hardness amplification* refers to methods that can be used to transform T into an even harder task, requiring even more resources. For example, one could consider the task of solving T on d inputs simultaneously. Is this d times harder, or can a clever reuse of resources allow for a computation that achieves more than what one would naively expect?

This question arises in many areas of computational theory, including cryptography, pseudorandomness, and optimization. Our foam construction is motivated by progress in understanding hardness amplification in the context of solving constraint satisfaction problems, a major topic in computer science, operations research, statistical physics, and information theory (e.g., Achlioptas et al.¹). A constraint satisfaction problem is specified by a collection of constraints on n variables. For the purpose of this discussion, we shall restrict ourselves to *bivariate* constraints, of the form $c(x, y)$, where each constraint c depends on only two of the n variables. A solution to the problem is a setting of the variables that maximizes the number of constraints that are satisfied. For example, the well-known graph coloring problem can be viewed as a bivariate constraint problem. Associate a single variable x_ν for each vertex ν of the given graph and each edge $\{u, v\}$

gives us the constraint $x_u \neq x_v$. If the variables are allowed to take values from $\{1, 2, 3\}$, then the constraints are all satisfiable, if and only the graph is 3-colorable.

One way to solve a constraint satisfaction problem is to try all settings of the variables and count the number of constraints that each setting satisfies, but this method is very inefficient, requiring exponential time. If $P \neq NP$, one cannot find a solution that satisfies *all* the constraints efficiently. A seminal hardness amplification result, the PCP Theorem,^{4,5} from 1992, improved this to show that if $P \neq NP$, then efficient algorithms cannot even find approximate solutions that satisfy *nearly all*, a $(1 - \epsilon_0)$ fraction of the constraints, when a solution satisfying all the constraints is known to exist. Here, ϵ_0 is a small positive constant. Indeed, the proof actually shows that it is hard to approximate even the fraction of constraints that are satisfiable.

Raz's celebrated Parallel Repetition Theorem²³ from 1995 dramatically strengthened this hardness of approximation result: he showed that for every $\epsilon > 0$, if $P \neq NP$, then efficient algorithms cannot guarantee solutions that satisfy very few, an ϵ fraction of the satisfiable constraints. Raz achieves this through a transformation of constraints called *parallel repetition* that we describe next. Suppose we are given a set of constraints on n variables. These constraints can be used to define new constraints on n^d variables as follows: Each new variable corresponds to a d -tuple of variables from the old problem. For every d -tuple of the original constraints c_1, c_2, \dots, c_d , we obtain a new constraint $C(X, Y) = c_1(x_1, y_1) \wedge c_2(x_2, y_2) \wedge \dots \wedge c_d(x_d, y_d)$ that can be thought of as a bivariate constraint on the variables $X = (x_1, \dots, x_d)$ and $Y = (y_1, \dots, y_d)$. The resulting compound constraints are said to have been obtained by repeating the original constraints d times in *parallel*.

If the best assignment can satisfy $(1 - \epsilon)$ fraction of the original constraints, how many constraints can be satisfied in the parallel repetition? Intuitively, one might think that the fraction of satisfiable constraints should be smaller, perhaps decay exponentially in d , since each new constraint corresponds to satisfying a set of d of the original constraints. But proving that this is true turned out to be very challenging, and counterexamples for pure exponential decay were known. In a breakthrough result, Raz showed that no assignment that can satisfy roughly $(1 - \epsilon^{32})^d$ fraction of the constraints obtained by parallel repetition, namely, *some* exponential decay, occurs. A key quantity of interest here is the *rate* of decay, namely, how "quickly" does the problem become harder, and the approximating factor decrease. Set $f(d)$ to be the largest number for which d repetitions of a constraint satisfaction problem decreases the approximation factor from $1 - \frac{1}{f(d)}$ to some small constant, say $1/10$. In this notation, Raz showed that $f(d) \geq d^{1/32}$, and left open determining the optimal dependence on d .

This result played a key role in showing that many types of problems cannot be solved by efficient algorithms (again assuming $P \neq NP$). A theorem that would prove a bound of the type $f(d) \approx d$ was dubbed a *Strong Parallel Repetition* theorem, and it remained open whether such a theorem holds. Subsequent works improved Raz's bound

toward such a strong theorem, first to $f(d) \geq d^{1/3}$ ¹⁵ in general, and then to $f(d) \geq \sqrt{d}$ ²² for an important subclass of constraints satisfaction problems, but the progress stopped at \sqrt{d} .

Feige et al.¹¹ were the first to observe that the parallel repetition question was related to foams. They studied a particular collection of constraints called *Odd Cycle* constraints, and showed that for them $f(d) \geq \sqrt{d}$. They also showed that if there is a d -dimensional cubical foam with surface area $A(d)$, then for this constraint satisfaction problem $f(d) \leq A(d)$. In particular, this meant that any proof that improved on their bound would show that there is no cubical foam with surface area \sqrt{d} , and any Strong Parallel Repetition theorem would prove that standard cubes are essentially the best cubical foams.

Once again, Raz²⁴ resolved the matter and showed that Odd Cycles were a counterexample to Strong Parallel Repetition. He proved that for Odd Cycle constraints $f(d) \leq \sqrt{d}$ and thus that the results of bounds $f(d) \geq \sqrt{d}$ of Feige et al.¹¹ and Rao²² are optimal! Indeed, one can view Raz's work as constructing a *discrete* cubical foam, with surface area proportional to \sqrt{d} . A key tool used by Raz was the so-called *Consistent Sampling Lemma*, invented by Holenstein to prove the upper bound of $f(d) \leq d^{1/3}$ mentioned above. Raz's result inspired our construction.

2. AN ALGORITHM FOR BUILDING CUBICAL FOAMS

Our solution to the Cubical Foam Problem involves generalizing Raz's discrete methods to Euclidean space and opening up the proof of the Consistent Sampling Lemma. We use the "Buffon's Needle" method to estimate surface area and we optimize our results using Fourier analysis.

Before describing our "sphere-like" cubical foam, we give some motivation for its construction. As stated earlier, our construction can also be interpreted as a very noise-resistant randomized discretization procedure for rounding off vectors of real numbers to vectors of integers.

DEFINITION 1. A discretization is a mapping which assigns each point $x = (x_1, \dots, x_d) \in \mathbb{R}^d$ to an integer point $r = (r_1, \dots, r_d) \in \mathbb{Z}^d$, such that $|r_i - x_i| < 1$ for each i . If a discretization has the property that whenever x is rounded to r , then $x + s$ is rounded to $r + s$ for all $s \in \mathbb{Z}^d$, we say that the discretization is periodic. Given a periodic discretization, we define its principle bubble to be the set of points that are rounded to the origin.

The principal bubble of a periodic discretization tiles \mathbb{R}^d according to the integer lattice. Thus, any periodic discretization immediately yields a cubical foam. We will in fact give a randomized procedure whose output is a periodic discretization (hence also a cubical foam). As described earlier, we say that such a procedure is noise-resistant if every two nearby points $x, y \in \mathbb{R}^d$ are unlikely to be assigned to different integer points. Intuitively, we expect the bubbles produced by a noise-resistant procedure to have small surface area, because x and y are assigned to different integer points only if the line segment joining them crosses the surface of a bubble. We will later see that finding a periodic discretization in which nearby pairs x and y are usually assigned to the same integer point is very similar

to a constraint satisfaction problem where every constraint involves two variables which are points in \mathbb{R}^d . Indeed, Raz's analysis of the Odd Cycle constraints suggests the investigation of a related problem (which we state somewhat imprecisely for the sake of brevity): Assign each point x in the unit cube $[0, 1]^d$ a shift $z \in [0, 1]^d$ in such a way that most nearby pairs of points are likely to be assigned the same shift.

The construction will be based on a carefully chosen probability distribution f on $[0, 1]^d$. For now, we describe the construction for any distribution f and later explain how to choose the optimal f . Let f_x be f translated by $-x$, so f_x is a probability density function on $-x + [0, 1]^d$; and let \tilde{f}_x be the periodic extension of this function, so $\tilde{f}_x(z) = f(z + x \bmod 1)$. Holenstein's Consistent Sampling Lemma gives a method for assigning shifts to all points such that the probability x and y are assigned different shifts is essentially $\int_{[0,1]^d} |f_x - f_y|$. We draw a sequence of points $(Z_1, H_1), (Z_2, H_2), \dots$ such that Z_i is drawn uniformly from $[0, 1]^d$, and H_i is uniform on $(0, \|f\|_\infty)$, where $\|f\|_\infty$ denotes the maximum value of f . Each x is then assigned a shift $z = Z_i$, where z_i comes from the first pair satisfying $\tilde{f}_x(Z_i) > H_i$. Thus, we are led to seek a function f for which this is small whenever x and y are nearby points.

Given a density function f and a number $h > 0$, consider the shape $D = \{x : f(x) > h\} \subset (0, 1)^d$ together with its boundary. We call D , or a translate of D a *droplet*. We will want droplets to have smooth boundaries, which do not touch the boundary of $[0, 1]^d$. For this reason, we will require that f 's periodic extension \tilde{f} be analytic and equal to 0 on the boundary of $[0, 1]^d$; we will call such a density function *f proper*. Given a proper density function f , we can now describe our randomized algorithm for producing a periodic discretization and associated cubical foam:

Algorithm 1: Periodic discretization (and foam) construction, given f :

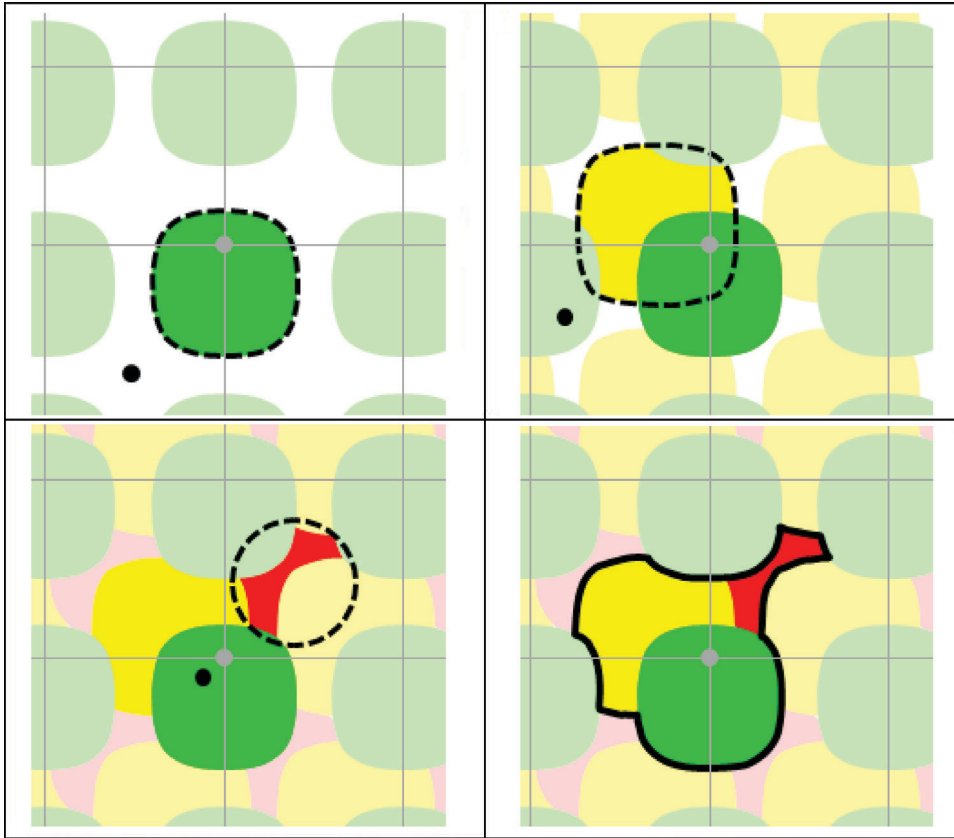
1. Let all points of \mathbb{R}^d be unassigned.
2. For stage $i = 1, 2, \dots$ until all points are assigned:
 - (a) Choose a uniformly random pair (Z_i, H_i) from $[0, 1]^d \times (0, \|f\|_\infty)$.
 - (b) Let droplet D_i be $\{x : f_{Z_i}(x) > H_i\}$, together with its boundary.
 - (c) Assign all currently unassigned points in D_i to $(0, \dots, 0)$, and extend this assignment periodically.
 - (d) Color all assignments from stage i with color i .

We remark that we color all the assignments merely to aid in the analysis of the algorithm.

It is not hard to prove that this algorithm indeed ends after a finite number of stages with probability 1. It is also clear that regardless of the algorithm's random choices, it always produces a periodic discretization. Thus, the points assigned to $(0, \dots, 0)$ by the algorithm always constitute a principal bubble, which partitions space according to the integer lattice.

We illustrate a sample run of the algorithm in Figure 3, with $d = 2$ and $f(x_1, x_2) = 4\sin^2(\pi x_1) \sin^2(\pi x_2)$. The integer lattice is outlined in gray, with the origin depicted as a gray dot. The first three panels illustrate stages 1, 2, and 3 of the

Figure 3. A sample run of the algorithm, with $d = 2$, $f(x_1, x_2) = 4\sin^2(\pi x_1) \sin^2(\pi x_2)$. The integer lattice is outlined in gray, with the origin depicted as a gray dot. The first three panels illustrate stages 1, 2, and 3 of the algorithm. In each stage, the black dot represents $-Z_i$ and the black dashed line outlines the droplet D_i . Colors 1, 2, and 3 are green, yellow, and red, respectively; we have used dark colors to show the points assigned to $(0, 0)$ and light colors to show their periodic translations. In the first panel is the assignment after stage 1: all points in the dark green droplet are assigned to $(0, 0)$; the light green translations are assigned periodically. The assignment after stage 2: the unassigned (uncolored) points within the outlined droplet are colored dark yellow and are assigned to $(0, 0)$. The assignment after stage 3, using red. The algorithm terminates after this stage—all points in \mathbb{R}^2 have been assigned. In the final panel, we outline the final bubble which partitions \mathbb{R}^2 periodically according to the integer lattice.



algorithm. In each stage, the black dot represents $-Z_i$ and the black dashed line outlines the droplet D_i . Colors 1, 2, and 3 are green, yellow, and red, respectively; we have used dark colors to show the points assigned to $(0, 0)$ and light colors to show their periodic translations.

3. ANALYSIS OF THE ALGORITHM

First, we compute the probability of rounding a pair of points x, y to different droplets:

THEOREM 1. *Let f be a proper density function, and \overline{xy} be a short line segment in \mathbb{R}^d , say $y = x + \epsilon \cdot u$, where u is a vector of length 1, and $\epsilon > 0$ is sufficiently small. For a given execution of Algorithm 1, let N denote the number of times \overline{xy} crosses the boundary between differently colored regions. Then*

$$\mathbb{E}[N] \approx \epsilon \cdot \int_{[0,1]^d} |\langle \nabla f, u \rangle|,$$

where the notation \approx means equality up to order ϵ^2 .

Next, using the relationship between noise resistance and surface area we show:

THEOREM 2. *Given an execution of Algorithm 1, let A denote the surface area of the boundary between color regions within $[0, 1]^d$. Then*

$$\mathbb{E}[A] = \int_{[0,1]^d} \|\nabla f\|.$$

Finally, we find f that minimizes the noise resistance and surface area:

THEOREM 3. *$f = \prod_{i=1}^d (2\sin^2(\pi x_i))$ is a proper density function with $\int \|\nabla f\| \leq 2\pi\sqrt{d}$. Moreover, for each vector of unit length u , f satisfies $|\langle \nabla f, u \rangle| \leq 2\pi$.*

The bounds we obtain for our cubical foam solution and for the noise resistance of our coordinated discretization procedure follow easily from these theorems. The bubble B output by Algorithm 1 has a surface area at most $2A$, where A is the quantity in Theorem 2; hence with the f from Theorem 3, the expected value of B 's surface area is at most $4\pi\sqrt{d}$. Hence, there must exist a bubble that tiles \mathbb{R}^d according to the integer lattice with a surface area at most $4\pi\sqrt{d}$. As for the noise resistance of Algorithm 1 as a coordinated

discretization procedure, if points $x, y \in \mathbb{R}^d$ at distance ϵ are assigned different integer points by the algorithm, then N , the number of times segment \overline{xy} crosses the boundary between color regions, must be at least 1. The probability of this event is at most $\mathbb{E}[N]$. Combining Theorems 1 and 3, this probability is at most $2\pi\epsilon$ (up to an error of ϵ^2 , but this error can be eliminated) as claimed.

Next we outline the proofs of Theorems 1, 2, and 3.

We begin with the main result, Theorem 1. Let $\tilde{D} = \text{closure}(\{x : \tilde{f}Z(x) > H\})$ be a random “droplet pattern” as would be chosen in a single stage of Algorithm 1; that is, \tilde{D} , is a random droplet along with all of its integer translations. Let I denote the event that \tilde{D} intersects the segment \overline{xy} , and let M denote the number of intersection points between the boundary of D and the segment \overline{xy} . Since \overline{xy} is very short, the first color region that touches the segment is very likely to completely enclose it. Thus, even *conditioned* on event I occurring, M is quite likely to be 0. More precisely, our main goal is to show that

$$\mathbb{E}[M|I] \approx \epsilon \cdot b, \quad \text{where } b = \int_{[0,1]^d} |\langle \nabla f, u \rangle|. \quad (1)$$

Using Equation (1), it is not hard to prove Theorem 1. To see this, consider the first stage of the algorithm’s execution in which a droplet pattern touching \overline{xy} is chosen. Let M_1 be the number of intersections between \overline{xy} and the boundary of this droplet pattern. Equation (1) tells us that $\mathbb{E}[M_1] \approx \epsilon \cdot b$. Recalling that N denotes the number of times \overline{xy} crosses the boundary between differently colored regions at the termination of the algorithm, we certainly have $N \geq M_1$. However, it is unlikely that N will exceed M_1 . Indeed, consider the second stage of the algorithm’s execution in which a droplet pattern touching \overline{xy} is chosen, and define M_2 to be the number of intersections between the boundary of this droplet pattern and \overline{xy} . Again, by Equation (1) we have $\mathbb{E}[M_2] \approx \epsilon \cdot b$. But furthermore, these M_2 intersections can only contribute to N if the first droplet pattern to touch \overline{xy} failed to completely enclose it. The probability of this is precisely $\Pr[M_1 > 0] \leq \mathbb{E}[M_1] \approx \epsilon \cdot b$. Hence, the expected contribution to N from this second stage is at most $(\epsilon \cdot b)^2$. Continuing the argument, we are able to upper-bound $\mathbb{E}[N]$ by

$$\epsilon \cdot b + (\epsilon \cdot b)^2 + (\epsilon \cdot b)^3 + \dots \approx \epsilon \cdot b.$$

We now describe the proof of Equation (1). Since the probabilistic choice of \tilde{D} is invariant under any translation of \mathbb{R}^d , we may assume $x = 0$ and hence $y = \epsilon \cdot u$. For a given $z \in [0, 1]^d$, let $g_z : [0, \epsilon] \rightarrow \mathbb{R}^0$ denote the restriction of \tilde{f}_z to the line segment from 0 to $\epsilon \cdot u$ and write $G(z) = \|g_z\|_\infty$. The event I occurs when the randomly chosen Z and H satisfy $H < G(Z)$, and the quantity M is equal to $\#\{s \in [0, \epsilon] : g_z(s) = H\}$. (Here, we have discounted events with probability 0.) Hence,

$$\mathbb{E}[M|I] = \frac{\int_{[0,1]^d} \int_0^{G(z)} \#\{s \in [0, \epsilon] : g_z(s) = H\} dH dZ}{\int_{[0,1]^d} G(Z) dZ} \quad (2)$$

Regarding the denominator of this expression, $G(Z) = gZ(0) = f(Z)$ up to an additive error of order ϵ by Taylor’s theorem,

and hence $\int_{[0,1]^d} G(Z) dZ = \int_{[0,1]^d} f(Z) dZ = 1$ up to order ϵ . As for the numerator of Equation (2), the inner integral equals the vertical distance traveled by a particle moving along the curve g_z ; this can be seen by partitioning the curve into small, nearly straight arcs, and considering their contribution to the integral. Hence, the inner integral equals $\int_0^\epsilon |g'_z(t)| dt$. Since $|g'_z(t)| = |g'_z(0)|$ up to order ϵ on $[0, \epsilon]$, $\int_0^\epsilon |g'_z| \approx \epsilon \cdot |g'_z(0)| = \epsilon \cdot |\langle \nabla \tilde{f}_z(0), u \rangle| = \epsilon \cdot |\langle \nabla f(Z), u \rangle|$ up to order ϵ^2 . Substituting this into the outer integral in the numerator of Equation (2), we conclude that

$$\mathbb{E}[M|I] \approx \epsilon \cdot \int_{[0,1]^d} |\langle \nabla f(Z), u \rangle| dZ = \epsilon \cdot b, \quad (3)$$

up to order ϵ^2 as claimed in Equation (1).

To prove Theorem 2, we need a method for computing surface area. We use the following (see Santalo²⁵)

THEOREM 4 (BUFFON’S NEEDLE THEOREM): *Let S be a \mathbb{Z}^d -periodic piecewise smooth surface. “Drop a needle” of length $0 < \epsilon < 1$, i.e., let x be a random point in $[0, 1]^d$, let u be a random vector of length 1, and let $y = x + \epsilon \cdot u$. If N denotes the number of intersections of the needle \overline{xy} with the surface S , then*

$$\text{area}(S \cap [0, 1]^d) = \frac{\mathbb{E}[N]}{c_d \cdot \epsilon}, \quad (4)$$

where c_d is the dimension-dependent-constant $\mathbb{E}[|u_1|]$.

Note that the value of c_d is not important, as it gets cancelled out in our analysis. Given the execution of Algorithm 1, we can apply Buffon’s Needle Theorem to compute the quantity A from Theorem 2. We obtain that $\mathbb{E}[A] = \mathbb{E}[N]/(c_d \cdot \epsilon)$, where in $\mathbb{E}[N]$ the probabilistic experiment is both the execution of Algorithm 1 and the random choice of the “needle.” Applying Theorem 1 for each choice of the needle, we obtain

$$\begin{aligned} \mathbb{E}[A] &= \mathbb{E}_u \left[\epsilon \cdot \int_{[0,1]^d} |\langle \nabla f, u \rangle| \right] / (c_d \cdot \epsilon) \\ &= \int_{[0,1]^d} \mathbb{E}_u [|\langle \nabla f, u \rangle|] / c_d, \end{aligned} \quad (5)$$

up to an additive error of order ϵ . Since ϵ can be arbitrarily small, Equation (5) is in fact an exact equality. And further, for each fixed vector $\nabla f(z)$, the quantity $\mathbb{E}_u [|\langle \nabla f(z), u \rangle|]$ equals $\|\nabla f(z)\| \cdot c_d$. Substituting this into Equation (5) proves Theorem 2.

It remains to prove Theorem 3. Suppose, we seek a proper density function on $[0, 1]^d$ such that $\int \|f\|$ is small. Writing $f = g^2$ and using the Cauchy–Schwarz inequality:

$$\int \|\nabla f\| = 2 \int |g| \cdot \|\nabla g\| \leq 2 \sqrt{\int g^2} \cdot \sqrt{\int \|\nabla g\|^2}, \quad (6)$$

where we also used that $\int f = \int g^2 = 1$. Since f is a proper density function, g ’s periodic extension is smooth and 0 on the boundary of $[0, 1]^d$. We may therefore rewrite g using the multidimensional sine series:

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

Differentiating Equation (7) and applying Parseval’s theorem, we get

$$\int \|\nabla g\|^2 = \pi^2 \sum_{\omega \in \mathbb{N}^d} \|\omega\|^2 \hat{g}(\omega)^2. \tag{8}$$

Applying Parseval to $\int g^2 = 1$ yields that $\sum_{\omega \in \mathbb{N}^d} \hat{g}(\omega)^2 = 1$, subject to which Equation (8) is minimized when $\hat{g}(1, 1, \dots, 1) = 1$. Thus, we are led to the solution for f stated in Theorem 3 and obtain the bound $\int \|\nabla f\| \leq 2\pi\sqrt{d}$ from Equations (6) and (8). It remains to verify that $\int |\langle \nabla f, u \rangle| \leq 2\pi$ also holds for each vector u . Using the Cauchy-Schwartz inequality, we obtain

$$\int |\langle \nabla f, u \rangle| \leq 2\sqrt{\int \langle \nabla g, u \rangle^2} = 2\sqrt{\sum_i u_i^2 \pi^2} = 2\pi,$$

for our choice of g , and hence f . This completes the proof of Theorem 3.

4. A THREE-DIMENSIONAL CUBICAL FOAM

Although we have asymptotically solved the Cubical Foam Problem up to a small constant factor, in the physically natural case of $d = 3$, our construction does not improve on the Choe Prism, or even the cube. Here, we present an improved three-dimensional cubical foam, constructed via an ad hoc method.

The two-dimensional minimizer given by Choe in Figure 2 (left) can be described as follows: Start with a

“base” facet centered at the origin; specifically, an edge from $(-s, s)$ to $(s, -s)$ for some parameter s . This already gives all vertices, by periodic extension. The hexagonal bubble is the convex hull of the two base points, their translates within $[0, 1]^2$, and their translates by $(1, 1)$. One chooses s to minimize the resulting surface area (perimeter).

We similarly construct a tiling shape B in three dimensions. We form a “base” facet centered at the origin, which is a regular hexagon, with vertices $\pm(0, -t, t)$, $\pm(-t, 0, t)$, and $\pm(-t, t, 0)$, for some $t \in (0, 1/3)$. Again, this gives all vertices, by periodic extension. We take B to be the convex hull of the 6 base points, along with their 6 translates within $[0, 1]^3$ and their 6 translates within $(0, 1]^3$. The polytope B has 14 facets: two opposing base regular hexagons, six larger “isosceles” hexagons, and six rectangles. An illustration is in Figure 4.


One may calculate that B has a surface area

$$6\left(\sqrt{3t^2 + \sqrt{2t}\sqrt{1-4t+6t^2}} + (1-t)\sqrt{1-2t+3t^2}\right),$$

which is minimized when $t \approx 0.1880$, having minimal value about 5.6121. This already beats the surface area of the Choe prism.

We can further improve this solution by letting B relax (within the torus $\mathbb{R}^3/\mathbb{Z}^3$) as a soap bubble. Using Brakke’s

Figure 4. Our new three-dimensional cubical foam. (Top left) The unrelaxed tile. (Top right) The tile after it has relaxed according to Plateau’s Rules. (Bottom left) The unrelaxed tile forming a foam according to the integer lattice. (Bottom right) Illustration of the relaxed foam as soap bubbles.



Surface Evolver,⁷ we obtain the relaxed bubble B shown in Figure 4. We remark that it has slightly wavy faces and curved edges, and that the vertices have moved according to $t \approx 0.1814$. The surface area of B is slightly less than 5.602, according to Surface Evolver.

5. DISCUSSION

We have given a probabilistic construction of a cubical foam with near-spherical surface area. The construction uses ideas that are new to the study of foams, and is inspired by work on the limits of “hardness amplification” in certain computational optimization problems. Our construction gives the first suggestion that in high dimensions, optimal foams might not be derived from Voronoi cells and may be quite unlike polyhedra.

We have also given an algorithmic application of our foam’s construction: a very “noise-resistant” procedure for rounding off vectors of d real numbers to integers. This discretization algorithm may not be practical for very large d , as Algorithm 1 is likely to run for a number of stages which is exponential in d . An important open problem is to find a coordinated discretization procedure with similar noise resistance, but taking time that grows only polynomially in d .


Finally, the construction of our cubical foam used randomness in an essential way; randomness is also used in other efficient high-dimensional constructions of foams (such as high-dimensional Kelvin foams). Although randomness is clearly required for noise-resistant coordinated discretization, it is an intriguing question as to whether it is necessary for the construction of foams, or whether explicit or derandomized constructions exist.

Subsequent to our work,¹⁷ Alon and Klartag² gave a simpler derivation of our cubical foam via Cheeger’s isoperimetric inequality; their analysis also shows that there exists a fixed parameter h that can be used as H_i throughout Algorithm 1. In other words, a good foam can be derived from the random translations of a single droplet of the form

$$D = \text{closure} \left(\left\{ x \in \mathbb{R}^d : \prod_{i=1}^d (2 \sin^2(\pi x_i)) > h \right\} \right).$$

However, it still remains unknown as to how to construct an explicit “spherical cube.”

Acknowledgments

We thank N. Alon, K. Ball, H. Cohn, P. D’Ancona, M. Goresky, J. Håstad, L. Hoffman, J. R. Lee, A. Klivans, D. Micciancio, V. Milman, O. Regev, and J. Sullivan for their insights. We thank R. Nelson for his assistance with raytracing code. Kindler was supported in part by the Koshland fellowship and by the Binational Science Foundation (BSF). O’Donnell was supported in part by NSF grants CCF-0747250 and CCF-0915893, BSF grant 2008477, and Sloan, Okawa, and von Neumann fellowships. Rao was supported in part by NSF grant CCF-1016565 and a Sloan fellowship. Wigderson was supported by NSF grants CCF-0832797 and DMS-0835373. 

References

- Achlioptas, D., Naor, A., Peres, Y. Rigorous location of phase transitions in hard optimization problems. *Nature* 435 (2005), 759–764.
- Alon, N., Klartag, B. Economical toric spines via Cheeger’s inequality. *I* (Sept. 18, 2009), 101–111.
- Arnold, W.N., Lacy, J.S. Permeability of the cell envelope and osmotic behavior in *Saccharomyces cerevisiae*. *J. Bacteriol.* 131, 2 (1977) 564–571.
- Arora, S., Lund, C., Motwani, R., Sudan, M., Szegedy, M. Proof verification and the hardness of approximation problems. *J. ACM* 45, 3 (1998), 501–555.
- Arora, S., Safra, S. Probabilistic checking of proofs: A new characterization of np. *J. ACM*, 45, 1 (1998), 70–122.
- Ball, P. Science in culture: Beijing bubbles. *Nature* 448, 7151 (2007), 256.
- Brakke, K.A. The surface evolver. *Exp. Math.* 1, 2 (1992), 141–165.
- Butler, G. Simultaneous packing and covering in Euclidean space. *Proc. London Math. Soc.* 3, 4 (1972), 721–735.
- Charikar, M., Chekuri, C., Goel, A., Guha, S., Plotkin, S. Approximating a finite metric by a small number of tree metrics. In *Proceedings of 39th Annual IEEE Symposium on Foundations of Computing* (1998), 379–388.
- Choe, J. On the existence and regularity of fundamental domains with least boundary area. *J. Differ. Geom.* 29 (1989), 623–663.
- Feige, U., Kindler, G., O’Donnell, R. Understanding parallel repetition requires understanding foams. In *Proceedings of 22nd Annual IEEE Conference on Computational Complexity* (2007), 179–192.
- Frank, F.C., Kasper, J.S. Complex alloy structures regarded as sphere packings. I. Definitions and basic principles. *Acta Cryst.* 11, 3 (1958), 184–190.
- Hales, T.C. The honeycomb conjecture. *Disc. Comp. Geom.* 25, 1 (2001), 1–22.
- Hales, T.C. A proof of the Kepler conjecture. *Ann. Math.* 162, 3 (2005), 1065–1186.
- Holenstein, T. Parallel repetition: simplifications and the no-signaling case. In *Proceedings of 39th ACM Symposium on Theory of Computing* (2007), 411–419.
- Holenstein, T. Parallel repetition: simplification and the no-signaling case. *Theory of Computing* 5, 1 (2009), 141–172.
- Kindler, G., O’Donnell, R., Rao, A., Wigderson, A. Spherical cubes and rounding in high dimensions. In *FOCS* (2008), IEEE Computer Society, 189–198.
- Kusner, R., Sullivan, J.M. Comparing the Weaire–Phelan equal-volume foam to Kelvin’s foam. *Forma* 11, 3 (1996), 233–242.
- Lee, J.R., Naor, A. Extending Lipschitz functions via random metric partitions. *Invent. Math.* 160, 1 (2005), 59–95.
- Nielsen, L.E., Landel, R.F. *Mechanical Properties of Polymers and Composites*, 2nd edn, CRC Press, Boca Raton, 1993.
- Plateau, J. *Statique expl’rimetale et thA’oretique des liquides soumis aux seules forces molA’culaires*, Gauthier-villars, Paris, 1873.
- Rao, A. Parallel repetition in projection games and a concentration bound. In *Proceedings of 40th ACM Symposium on Theory of Computing* (2008), 1–10.
- Raz, R. A parallel repetition theorem. *SIAM J. Comput.* 27, 3 (1998), 763–803.
- Raz, R. A Counterexample to Strong Parallel Repetition. *IEEE Computer Society*, 2008, 369–373.
- Santaló, L. *Integral Geometry and Geometric Probability*, 2nd edn, Cambridge University Press, Cambridge, 2002.
- Thomson, W. On the division of space with minimum partitional area. *Philos. Mag.* 24 (1887), 503–514.
- Weaire, D., Phelan, R. A counterexample to Kelvin’s conjecture on minimal surfaces. *Phil. Mag. Lett.* 69, 2 (1994), 107–110.

Guy Kindler (guy.kindler@weizmann.ac.il), School of Computer Science and Engineering, Hebrew University of Jerusalem.

Ryan O’Donnell (odonnell@cs.cmu.edu), Computer Science Department, School of Computer Science, Carnegie Mellon University.

Anup Rao (anuprao@cs.washington.edu), Computer Science and Engineering, University of Washington.

Avi Wigderson (avi@math.ias.edu), School of Mathematics, Institute for Advanced Study.