

THE GEOMETRY OF BUBBLES AND FOAMS

JOHN M. SULLIVAN

University of Illinois

Department of Mathematics

Urbana, IL, USA 61801-2975

Abstract. We consider mathematical models of bubbles, foams and froths, as collections of surfaces which minimize area under volume constraints. The resulting surfaces have constant mean curvature and an invariant notion of equilibrium forces. The possible singularities are described by Plateau's rules; this means that combinatorially a foam is dual to some triangulation of space. We examine certain restrictions on the combinatorics of triangulations and some useful ways to construct triangulations. Finally, we examine particular structures, like the family of tetrahedrally close-packed structures. These include the one used by Weaire and Phelan in their counterexample to the Kelvin conjecture, and they all seem useful for generating good equal-volume foams.

1. Introduction

This survey records, and expands on, the material presented in the author's series of two lectures on "The Mathematics of Soap Films" at the NATO School on Foams (Cargese, May 1997). The first lecture discussed the differential geometry of constant mean curvature surfaces, while the second covered the combinatorics of foams.

Soap films, bubble clusters, and foams and froths can be modeled mathematically as collections of surfaces which minimize their surface area subject to volume constraints. Each surface in such a solution has constant mean curvature, and is thus called a CMC surface. In Section 2 we examine a more general class of variational problems for surfaces, and then concentrate on CMC surfaces. Section 3 describes the balancing of forces within a CMC surface. Then in Section 4 we look at existence results for bubble

clusters, and at the singularities in such clusters and in foams; these are described by Plateau's rules.

These rules for singularities imply that a foam is combinatorially dual to some triangulation of space, as explained in Section 5. Triangulations of a surface are related through their Euler number to the topology of the surface. In three dimensions, not as much can be said along these lines, but we describe certain results in Section 6. Among the most interesting triangulations, from the point of view of equal-volume foams, are those of the tetrahedrally close-packed (TCP) structures from chemistry, described in Section 7. The TCP structures observed in nature fit into an infinite mathematical class; constructions for some infinite families are described in Section 8. Finally, in Section 9, we discuss the application of these ideas to Kelvin's problem of partitioning space.

The author is partially supported by NSF grant DMS-9727859, and would like to thank Karsten Große-Brauckmann, Andy Kraynik, Rob Kusner, Frank Morgan, Mike O'Keefe and Nick Rivier for helpful conversations on foams and related structures, and J.-F. Sadoc for the invitation to speak at the NATO school in Cargese. Some of the figures here were originally created by the author to illustrate [26]. Fig. 2 was created by Große-Brauckmann with software by Bernd Oberknapp.

2. Variational problems for surfaces

We want to examine the geometry of a surface which (like those in foams) minimizes its surface area, given constraints on the enclosed volume. Remember that a surface in space has (at each point) two *principal curvatures* k_1 and k_2 ; these are the minimum and maximum (over different tangent directions) of the normal curvature. Because there is no way to globally distinguish the two, only symmetric functions of the k_1 and k_2 are physically meaningful. The *mean curvature* is their sum $H = k_1 + k_2$ (twice the average normal curvature), while the *Gauß curvature* is their product, $K = k_1 k_2$. The importance of K is that it is an *intrinsic* notion, meaning that it doesn't change if we bend a surface without stretching it (like rolling up a flat piece of paper). But our problems depend explicitly on the way the surface sits in space—its extrinsic geometry—and thus we will find the mean curvature H more useful.

Instead of minimizing the area of a surface, we might consider minimizing a more general energy for surfaces which could include curvature effects. The most natural candidate would be an integral (over the surface) of symmetric quadratic terms in the principal curvatures k_1 and k_2 . Such

an energy E for a surface M can be written as

$$E(M) = \int_M \left(a + b(H - H_0)^2 + cK \right) dA,$$

where a , b , c and H_0 are constants and dA is the area form on the surface.

By the Gauß-Bonnet theorem, however, the integral of K over M is just a topological constant: it is $2\pi\chi$ on a closed surface of Euler characteristic χ , and, for a surface with boundary, depends additionally just on the boundary data. Thus the final term in E stays constant as we vary M continuously, so it is uninteresting for our variational problem. Furthermore, if the two sides of the surface M are equivalent, then H_0 must be zero, and we are left with just two terms in E : a constant times the area of M , plus another times the so-called *Willmore energy* $W(M) := \int H^2 dA$. This scale-invariant elastic bending energy W for surfaces [55, 19] has been proposed as a good model for the geometry of certain biological membranes like lipid vesicles [30, 56, 20]. A surface which minimizes W subject to constraints fixing both area and volume, for instance, can look quite like the usual shape of a red blood cell.

In foams, however, the interfaces behave like soap films. Here, the curvatures don't seem to affect the energy, so the only term we will consider is surface tension (trying to minimize the area), perhaps balanced by a pressure differential (fixing the volume to one side of the surface). A perturbation of a surface M is given by a vector-valued function \vec{v} on M , telling how each point is to move. Let us consider the first variations of area A and volume V under this perturbation. If \vec{n} is the unit normal to M and $\vec{H} = H\vec{n}$ is the mean curvature vector, then

$$\delta V = \int_M \vec{v} \cdot \vec{n} dA, \quad \delta A = \int_M \vec{v} \cdot \vec{H} dA,$$

so we say that \vec{n} is the gradient of volume and \vec{H} is the gradient of area. If we write $H = \delta A / \delta V$, we obtain a second, variational, definition of mean curvature.

These Euler-Lagrange equations show that a surface is critical for the area functional exactly if $H \equiv 0$. Such a surface is called a *minimal surface*, and indeed any small patch of M is then area-minimizing for its boundary. On a minimal surface, $k_1 = -k_2$, so necessarily $K \leq 0$. Thus we see that a minimal surface is saddle shaped, as in Fig. 1. Thus it is never convex, meaning that a plane brought towards a patch of a minimal surface must touch first at the boundary of the patch, not in the interior. More generally, the (very useful) *maximum principle* says that the same happens whenever any two minimal surfaces are brought together (so that one is never even “relatively convex” compared to the other).

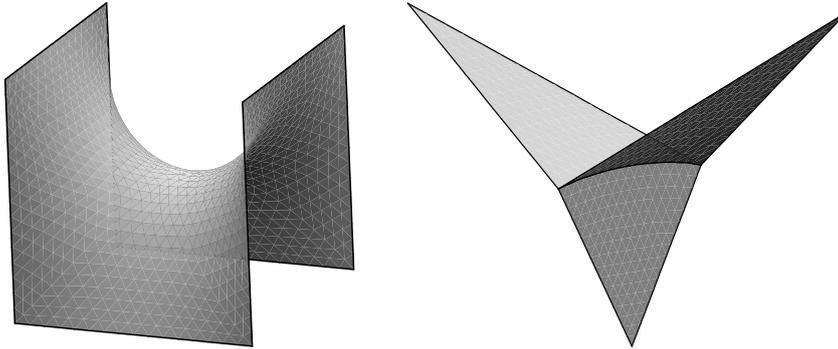


Figure 1. On the left, the least-area surface spanning the fixed edges of a cube. We see that, as in any minimal surface, each point is saddle shaped, with equal and opposite principal curvatures, and thus has negative Gaussian curvature. At the right, three minimal surfaces meet along a triple line at 120° . This surface would be achieved by soap film spanning three wire Vs, and when extended by two-fold rotation about these lines it forms the Kelvin foam.

If we realize that $\vec{H} = \Delta_M \vec{n}$, then we can think of minimal surfaces as a nonlinear version of harmonic functions. There is an extensive classical theory of minimal surfaces; good introductions can be found in [38, 10, 17] for instance. Much of the theory involves complete embedded minimal surfaces [18], and is based on the fact that the Gauß map $\vec{n}: M \rightarrow S^2$ is conformal; this leads to the Weierstraß representation, which parameterizes M in terms of holomorphic data on a Riemann surface (and can be reinterpreted in terms of spinors [25]).

A soap film spanning a wire boundary, with no pressure difference, will be a minimal surface, but in general we need to consider the effect of a volume constraint. In this case, we find that the surface has constant mean curvature $H \equiv c$, so it is called a *CMC surface*. Here the constant c is the Lagrange multiplier for the volume constraint, which is exactly the pressure difference across the surface; this relation $H = \Delta p$ is known as Laplace's law.

We have seen that a CMC surface is a critical point for area A , given the constraint of fixed volume V , with H being the Lagrange multiplier. Equivalently, if we know the pressure difference H , the surface is a critical point for $A - HV$. In physical bubble clusters or foams, each cell contains a certain fixed amount of air or other gas (neglecting slow diffusion effects). Of course, air is not an incompressible fluid. But for our mathematical theory, we will usually pretend it is, fixing a volume (rather than a mass of gas) in each cell. The point is that no matter what the relation between pressure and volume is for the real gas, the equilibrium structures will be the same for either problem.

For instance, a single bubble will adopt the form of a round sphere. This *isoperimetric theorem*, that a sphere minimizes surface area for its enclosed volume, was known to the ancient Greeks. (Indeed, a famous story credits Queen Dido with making use of the two-dimensional version when founding the city of Carthage.) But a proper mathematical proof (which must demonstrate the existence of a minimizer rather than merely describing properties of one if it does exist) was only given in the last century by Schwarz [44], using a symmetrization argument.

Note that the unit sphere (with $H = 2$) is an unstable critical point for the functional $A - 2V$, even though it is the unique minimizer for area A , given volume $V = 4\pi/3$. Although these problems have the same critical points, as discussed above, the stability of a critical point may differ.

3. Forces in CMC Surfaces

One of the most important ideas in the mathematical study of CMC surfaces is that of *force balancing*, as described in the work of Korevaar, Kusner and Solomon [21]. In general, Noether's theorem says that symmetries of an equation lead to conserved quantities. Here, the equation $H \equiv c$ for a CMC surface is invariant under Euclidean motions (translations and rotations) and the corresponding conserved quantities on a CMC surface can be interpreted as forces and torques.

The physical interpretation is as follows. Start with a complete embedded CMC surface M , and cut it along a closed curve γ . The two halves of M pull on each other with a force due to surface tension and to pressure (exerted across an arbitrary disk D spanning γ). If \vec{n} is the normal to D , and $\vec{\eta}$ is the conormal to γ in M (that is, the unit vector tangent to M but perpendicular to γ), then the force can be written as

$$\vec{F}(\gamma) = \int_{\gamma} \vec{\eta} - H \int_D \vec{n}.$$

Here we think of pressure $p = 0$ outside M , and $p = H$ inside M . Torques around any origin can be defined in a similar fashion.

By Stokes' theorem, \vec{F} is independent of the choice of D . It is also somewhat independent of γ , in the following sense. The fact that M is in equilibrium, with surface tension balancing pressure differences, means that no compact piece feels any net force or torque. If γ and γ' are homologous curves on M , their difference bounds such a compact piece, so their forces must be equal. That is, $\vec{F}(\gamma)$ actually depends only on the homology class of γ .

The CMC surfaces of revolution, called *unduloids*, were first studied by Delaunay [9] in the last century. An unduloid is a wavy cylinder as in

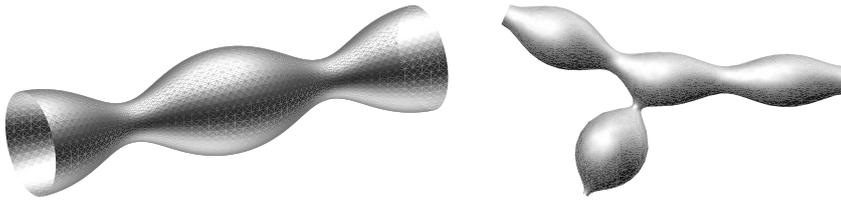


Figure 2. Left, a typical unduloid (a CMC surface of revolution), whose force vector points along the axis. Right, a typical triunduloid (embedded CMC surface with three ends). Each end is asymptotic to an unduloid; the three force vectors sum to zero.

Fig. 2; the force vector, for a (homologically nontrivial) curve γ around the cylinder, necessarily points along the axis of revolution. The family of unduloids (for given H) can be parameterized by the strength of this force, which is maximum for a straight cylinder. In fact, Delaunay described the generating curves for unduloids as solutions of a second-order ODE, and the force is a first integral for this equation. As the force decreases from that of a cylinder, we see alternating necks and bulges. The necks get smaller and smaller, like shrinking catenoids, and in the limit of zero force, the unduloid degenerates into a chain of spheres, as in a beaded necklace.

The idea of force balancing has been very useful in the classification of complete embedded CMC surfaces. The only compact example is the round sphere [1], with all forces equal to zero. In general, each end of such a surface is asymptotic to an unduloid [21]. Thus it has a force vector associated to it, directed along the axis of the unduloid, pulling the “center” of the surface out towards infinity. The conservation of forces (the homology invariance of \vec{F}) means that these forces on the ends must sum to zero. This *balancing condition* is the essential ingredient in classifying such CMC surfaces, and combined with some spherical trigonometry can lead to a complete classification of the three-ended surfaces [13, 14].

Of course, such infinite surfaces do not exist in nature; bubble clusters and foams are collections of CMC surfaces meeting along singular lines as described below in Section 4. We note that, if these singular lines were known in advance, each piece of surface spanning them could be viewed as a solution of the CMC Plateau problem [16] on its boundary. The ideas of forces and torques carry over to these bubble clusters and foams. We might most naturally cut a cluster M with a plane P (of normal \vec{n}); the force exerted across this cut is

$$\int_{P \cap M} \vec{\eta} - \int_{P \setminus M} p \vec{n},$$

where p is the pressure in each region of the complement of the surface M (again normalized to have $p = 0$ outside the cluster). For a finite cluster of

bubbles, these forces are always zero, since they are again constant as we move the plane P , and are evidently zero once P fails to hit M .

In an infinite periodic foam, however, there can be non-zero forces, and we can use them to calculate the *stress tensor* for the foam. This stress tensor σ shows how the lattice of periodicity for a periodic foam would like to change (by stretching or shearing), and is usually computed (as in [40] or [23]) as

$$\sigma = \int_{M_0} I - \vec{\nu}\vec{\nu}^T,$$

plus a multiple of the identity matrix which depends on the choice of absolute zero pressure. Here M_0 is one fundamental domain for the periodic foam, and ν at any point is the normal vector to M_0 . We claim that this tensor can also be computed in terms of the forces. Essentially, given any vector \vec{n} , the force when M_0 is cut by a plane with normal n gives $\sigma(\vec{n})$. We will use these force or stress calculations later to help minimize the total area of an equal-volume periodic foam of given combinatorics, by adjusting its lattice parameters.

4. Bubble Clusters and Singularities in Soap Films

So far, we have been rather cavalier about assuming that solutions to our problem (of minimizing surface area with constraints) exist and are smooth surfaces (hence of constant mean curvature). Of course, the surfaces in a foam or bubble cluster are only piecewise smooth, and the pieces meet in the singularities observed by Plateau [39].

Almgren proposed the following *bubble cluster problem*: given volumes V_1, \dots, V_k , enclose and separate regions in space with these volumes, using the least total area. He then proved that a minimizer exists and is a smooth surface almost everywhere [2] (see also [31]).

The existence proof uses a direct approach which takes the limit of a sequence of surfaces whose area approaches the infimal value. The difficulty here (and one reason the theorem is hard even for the case of a single region) is that surfaces of decreasing area can have growing “filigree”. Perhaps the n th surface M_n looks like the minimizer, but with a (bent or even branched) tube of length n and radius $1/n$ attached. (There is a nice figure in [32, p. 5].) The areas of these surfaces decrease to the proper minimum value, but every point in space might be a limit of the filigree. To prove the theorem, Almgren used a more sophisticated notion of limit surface from geometric measure theory, which automatically discards the filigree.

This suggests that the proof does not give any control on the topology of the regions; in fact we are not allowed even to assume that each region is connected. For all we know, it may be that the best way to make a

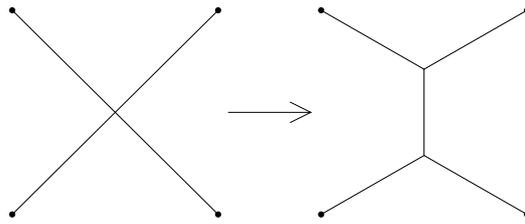


Figure 3. The X junction at the left is not an (M, ε, δ) -minimizing set in the plane because it can deform (at an arbitrarily small scale) to the double Y at the right. In the plane, the Y is the only singularity allowed.

cluster with some four fixed volumes is to split one of them into two pieces (with the correct total volume), and make what looks like a cluster of five bubbles. A reasonable conjecture, of course, is that in any minimizer, each region is connected.

Almgren's result does, however, say something about the shape of the minimizers: they are (M, ε, δ) -minimizing sets. This means, by definition, that a Lipschitz deformation of space in a small ball (of radius $r < \delta$) can reduce the area of the set at most by a factor of $1 + \varepsilon(r)$ (where $\varepsilon \rightarrow 0$ as $r \rightarrow 0$). If there are no volume constraints, a soap film, even with singularities, is $(M, 0, \delta)$ -minimizing. The function $\varepsilon(r)$ is a clever way to allow for non-zero mean curvature while not expanding the allowed set of singularities for soap films.

In the plane, an (M, ε, δ) -minimizing set is a finite network of smooth curves meeting in threes at equal 120° angles. Note that the deformation (shown in Fig. 3) of an X into a double Y (which is the Steiner tree on the four vertices of a square) shows that the X is not minimizing—it can be improved at any small scale. (Of course, the models we use here are meant to describe dry foams; in a wet foam with positive liquid fraction $\varphi > 0$, there are Plateau borders near each singularity. In this case one might wonder if additional singularities might be possible. Although a wet X is stable for large φ , Brakke and Morgan [4] have shown that it is unstable for small positive φ . The analogous question in three dimensions is still open.)

In space, Taylor was able to show [49] that the singularities observed empirically in soap films by Plateau are the only ones possible in an (M, ε, δ) -minimizing set. That is, such a minimizer consists of a finite union of smooth surfaces, meeting in threes at 120° dihedral angles along a finite number of smooth curves; these curves in turn meet at finitely many tetrahedral corners (which look like the soap film obtained when dipping a tetrahedral frame into soapy water—six sheets come together along four triple curves into the central singularity). The proof relies on ruling out seven other candidate cones, like that on the twelve edges of a cube; if we dip a

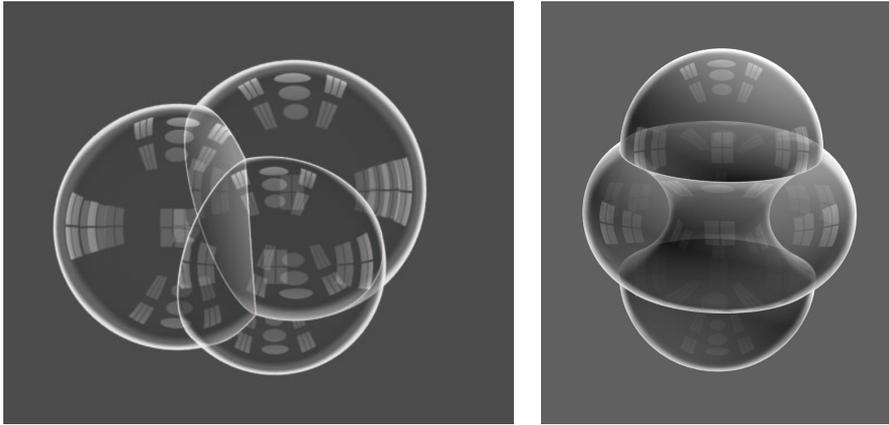


Figure 4. At the left we see the standard cluster of three soap bubbles, with all interfaces spherical, presumably minimizing total area for its enclosed volumes. At the right we see a possible competitor for the least-area double bubble, where one region is a toroidal belt around the other.

wireframe cube into soap, the film formed breaks the symmetry, inserting a small square at the center whose four vertices are all tetrahedral corners. Similarly, the other candidate cones all break apart into configurations of less area with several tetrahedral corners.

We thus know that the minimizing bubble clusters guaranteed by Almgren's theorem follow *Plateau's rules*, having just these possible singularities, as seen in the standard triple bubble in Fig. 4, left. Actually, we know more: here each smooth piece has constant mean curvature, and the mean curvatures (being pressure differences) satisfy a *cocycle condition*. This means that if we add up the mean curvatures of each surface we cross, as we follow a closed loop in space, the total is zero. Equivalently, numbers (call them pressures) can be assigned to each component of the complement such that the mean curvature of each interface is the difference of these pressures. Although we cannot guarantee that each region (of fixed volume) is a single component, the different components making up a single region will necessarily have the same pressure.

Of course a minimizing cluster of just one region will be a round sphere. It is surprising, though, how little is known about other small clusters. The case of two regions, a *double bubble*, is one of the geometrical problems where "many mathematicians believe, and all physicists know," that the obvious solution is correct.¹ The best double bubble should be the one seen when playing with soap bubbles, with three spherical caps meeting along a single circular triple junction. Surprisingly, this is not yet a theorem. However,

¹The quote is from Rogers [42], describing Kepler's sphere packing problem.

Hutchings showed (following a suggestion of White) that the minimizer must have rotational symmetry, and gave further conditions, which, in the case of equal volumes left only one family of competitors, as in Fig. 4 (right), allowing Hass and Schlafly to rule them out with a computer search [15] (see also [33]).

For other small clusters, the author has suggested [48] that the minimizer for up to four regions in any volume ratios should be a *standard cluster*, obtained by stereographic projection of a simplex in S^3 ; here the interfaces are all pieces of spheres and the triple junctions are all circular arcs. Perhaps even all stable clusters of up to five bubbles are spherical in this sense, but there is evidently a non-spherical cluster of six bubbles (shown in the author's illustration on the cover of [32]). Many more conjectures along these lines can be found in [48].

Almgren's theorem holds for finite clusters of soap bubbles in Euclidean space or certain other symmetric manifolds. Mathematically we would like to consider a foam as an infinite cluster. If we ask that the foam be periodic, then it is equivalently a finite cluster of bubbles in a torus (\mathbb{R}^3 mod the lattice of periodicity), so Almgren's theory does apply, and we get a minimizer. But for the general case of infinite foams (even with equal-volume cells) very little is known, though Morgan has made some progress on the two-dimensional case [34].

5. Combinatorics of Foams

We will define a foam mathematically as a (locally finite) collection of CMC surfaces, meeting according to Plateau's rules, and satisfying the cocycle condition: pressures can be assigned to each component of the complement so that the mean curvature of each interface is the pressure difference. We will call the components of the complement the *cells* of the foam, call the interfaces between them simply the *faces*, call the triple junction lines where they meet the (Plateau) *borders*, and call the tetrahedral singularities the *corners*.

As we have mentioned, finite clusters minimizing area for fixed volumes are known to have this structure. Although there are no analogous results about infinite clusters (except when they are periodic) this is a reasonable definition for foams, and we expect it to describe all physical foams whose shape is determined by surface tension.

For many purposes, we can ignore the geometric parts of Plateau's rules, and only pay attention to the combinatorial aspects of the foam's cell complex. The combinatorial rules mean that the *dual cell complex* to a foam is in fact a simplicial complex, that is, a triangulation of space. To construct this dual, we put a *vertex* in each cell of the foam. Vertices in a pair of

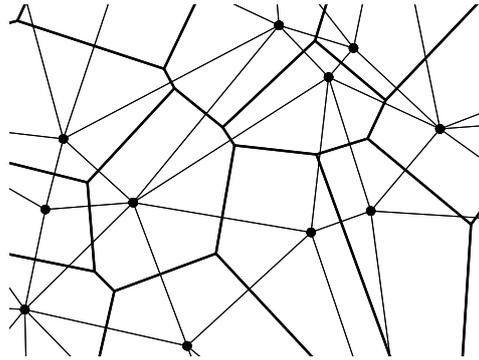


Figure 5. The heavy lines show a sample Voronoi partition in the plane, for the sites marked with dots. The thin lines are the edges of the dual Delone triangulation, which connect sites with adjacent Voronoi cells.

adjacent cells are connected by an *edge*, which is dual to a face of the foam. Where three faces come together along a border, we span the three corresponding edges with a *triangle*; and where four borders come together at a corner, we fill in a *tetrahedron*.

It is possible that in a foam, a cell or face might fail to be simply connected; or perhaps some face might have the same region on both sides or might meet itself along a triple junction. To avoid such technical complications, we will henceforth deal only with *simple foams*, where each cell (whether open, including only the interior, or closed, including the boundary) is a ball, and each (open or closed) face is a disk, and each border is an interval. Dually, these conditions mean, for instance, that no two tetrahedra intersect in more than one triangle; these conditions are often part of the definition of a triangulation. In any case, the foams we care about are all simple, having cells which approximate convex polyhedra.

Combinatorially, the Plateau rules mean that a (simple) foam and its dual triangulation are like a Voronoi decomposition of space and the dual Delone triangulation. Given a set of *sites* in space, the Voronoi cell [45, 35] for each site is the convex polyhedron consisting of points in space closer to that site than to any other. If the sites are in generic position (with no five on a common sphere) then the dual Delone complex (whose vertices are the original sites) is completely triangulated; each Delone tetrahedron has the property that no other sites are inside its circumscribed sphere. An example in the plane is shown in Fig. 5.

This similarity suggests that that we might look for foams as relaxations of Voronoi decompositions. For instance, we can give a modern interpretation to Kelvin's construction [50] of his candidate for a least-area partition of space into equal volume cells as follows. Start with sites in the body-

centered cubic (BCC) lattice. Their Voronoi cells are truncated octahedra, packed to fill space. If we let the films in this packing relax (until the geometric parts as well as the combinatorial parts of Plateau's rules are satisfied) we should get a periodic foam, the one proposed by Kelvin.

Starting from other periodic arrays of sites, we can carry out the same procedure to generate a foam. Mathematically, there is no theory for the relaxation step, which would follow a mean-curvature flow for surfaces with triple junctions and volume constraints. But in practice this procedure is a useful way to compute foam geometries. The author's `vcs` software for computing Voronoi diagrams [47] in three dimensions can send its output to Brakke's Surface Evolver [5]. This latter program can interactively manipulate, refine and evolve triangulated surfaces into energy-minimizing configurations, modeling the physical effects of surface tension and pressure. New symmetry features in the Evolver [6] allow computation of just one fundamental domain in a foam of arbitrary symmetry, thus giving greater accuracy for given computational effort.

6. Bounds on the combinatorics of triangulations

Consider a triangulation of any compact three-dimensional manifold, and let V , E , F and T be the numbers of vertices, edges, triangles and tetrahedra, respectively. Let z be the average number of edges at a vertex (which is the average number of faces on a cell of the dual foam), and n be the average valence of an edge (the average degree—or number of borders—of a face in the foam). Then a multiple-counting argument shows that

$$4T = 2F, \quad 3F = nE, \quad 2E = zV.$$

The Euler characteristic χ of a triangulation is the alternating sum of the number of simplices in different dimensions. For a triangulation of a surface, $\chi = V - E + F$, while multiple counting shows $3F = 2E = zV$. It follows that $6/z = 1 + 2\chi/F$, meaning that the sign of χ determines the sign of $6 - z$: triangulations of a sphere have $z < 6$, those of a torus have $z = 6$, and those of a surface of higher genus have $z > 6$.

Returning, however, to our three-dimensional triangulation, we know that the Euler characteristic of any compact 3-manifold is zero: $0 = \chi = V - E + F - T$. (In general, the Euler characteristic distinguishes different manifolds only in even dimensions.) Combining this with the multiple-counting results gives

$$6 - n = 12/z. \quad (*)$$

Unlike the case of surfaces, where z was related via χ to the topology of the manifold, equation (*) relates z to n independent of the ambient manifold. We have been talking about triangulations of compact manifolds, but a

periodic triangulation of an infinite manifold (like Euclidean or hyperbolic space) covers a triangulation of the compact quotient manifold, so it will satisfy the same relation (*).

For the regular triangulations of spherical space S^3 , n and z will be integers: $(n, z) = (3, 4)$ or $(4, 6)$ or $(5, 12)$. We can consider the regular ideal triangulation of hyperbolic space as having $(n, z) = (6, \infty)$. Of course, a regular Euclidean tetrahedron cannot tile flat space; its dihedral angle is $\arccos \frac{1}{3} \approx 70.53^\circ$, so we would need

$$n = n_0 := \frac{2\pi}{\arccos \frac{1}{3}} \approx 5.1043$$

in such a tiling. This would correspond by (*) to

$$z = z_0 := 12/(6 - n_0) \approx 13.397$$

which is 4π divided by the measure of the solid angle at the vertex of a regular tetrahedron.

A regular tetrahedron has dihedral angle $\arccos \frac{1}{3}$, and if we view average dihedral angle as a function on the space of tetrahedra, by symmetry it has a critical point at this regular tetrahedron (as the Shoemakers [46] have pointed out—see Section 7). Thus, if we tile Euclidean space by nearly regular tetrahedra, we expect the average dihedral angle to be very close to $\arccos \frac{1}{3}$. But the geometric part of Plateau's rules says that each corner in a foam has exactly tetrahedral angles. So, if the borders in the foam are not too curved, the dual triangulation should consist of nearly regular tetrahedra. Thus we expect $n \approx n_0$, so also $z \approx z_0$ by (*).²

Another way to think about this is to consider a face in a foam. If the face is nearly flat, and its borders are nearly straight, then its equal angles of $\arccos(-1/3) \approx 109.5^\circ$ suggest that it should be an n_0 -gon. Or, rather, that perhaps pentagons (and hexagons) might be favored in a foam. This has led many people to conjecture [41, 54, 8] that efficient foams should have many pentagonal faces, unlike the ordered Kelvin foam, but more like the disordered equal-volume foams observed in nature [28].

In the case of a foam with equal-pressure cells, all the interfaces are minimal surfaces. Here, Kusner [24] has given a clever argument to show that we must have $z \geq z_0$ and $n \geq n_0$, by applying the Gauß-Bonnet theorem to the boundary of each cell in the foam. The corners of the cell each contribute a known quantum of Gauß curvature (equal to $4\pi/(2z_0 - 4)$) because of their tetrahedral angles. The borders contribute an unknown

²Note that Coxeter [8] outlines a different argument (involving Petrie polygons) which suggests that $n \approx (13 + \sqrt{313})/6 \approx 5.115$ and hence $z \approx 13.56$.

amount of curvature to each of the three cells they bound, but the equal 120° dihedral angles along each border imply that these amounts sum to zero, and can be neglected in the average. Finally, the faces, being minimal surfaces, contribute a negative amount of curvature. We see that the average cell (having total curvature 4π) must thus have at least $2z_0 - 4$ corners, or equivalently at least z_0 faces.

Choe showed that, given any compact manifold, in its universal cover it has a fundamental domain of least boundary area [7]. For us, this means that if we look at any of the ten crystallographic groups that act freely on \mathbb{R}^3 (see for instance [51, §4.3]), there is a least-area foam with a single type of cell, tiling space by that group action. For instance, if the group contains merely translations, the quotient manifold is a torus, and, no matter which lattice parameters we choose, this least-area Choe foam presumably is combinatorially equivalent to Kelvin's foam. In particular, if we choose the BCC lattice, the conjecture [48] is that this foam is exactly Kelvin's foam. It seems likely also that the Williams foam [54] is the Choe foam for its symmetry group (which involves one screw motion with a 90° twist). For any of the ten groups (including four with nonorientable quotients) Choe's construction gives a foam with congruent equal-pressure cells, so by Kusner's result, each cell must have at least 14 faces. We know of no other examples of equal-pressure foams.

From the examples of regular triangulations, it is tempting to think that somehow, triangulations with higher values of z (or, equivalently, n) should tile hyperbolic manifolds, while those with lower values should tile spherical manifolds. Of course, we know that the Euler number can give no such information in three dimensions, but Luo and Stong [27] have given one result in this direction. They show that any triangulation with $z < 8$ ($n < 9/2$) must be a triangulation of (some quotient of) S^3 . And if $z = 8$, the only other possibility is (some quotient of) $S^2 \times \mathbb{R}$.

However, this seems to be the only topological information that can be extracted from z , because simple refinement constructions change the value of z for a triangulation. We can apply these iteratively (as shown in [27]) to get a sequence of triangulations with arbitrary $z > 8$ in the limit, on any three-manifold.

One construction refines a tetrahedron by adding a vertex in its interior, replacing it by four tetrahedra. (Dually, we blow a small tetrahedral bubble at one corner in the foam.) This replaces z by a value closer to 8. For instance, O'Keefe [36] has noted that if we start with the Kelvin foam (or sodalite structure) and decorate each corner with a new tetrahedral bubble,³ we change z from 14 to $62/7$, already close to 8.

³O'Keefe also proposes packing spheres at the corners of this "decorated sodalite" to give a sphere packing of least density.

That first construction suffices to reduce z . A second construction refines an edge of the triangulation, by adding a vertex at its midpoint. (Dually we inflate one facet of the foam to form a new prismatic bubble.) This construction applied repeatedly to a k -fold edge brings z close to $2 + 2k$. Since an edge of high degree can be created in any triangulation by the first construction, we can then achieve arbitrarily high z values this way.

As we have seen, with no restrictions on the triangulation, a value of $z > 8$ tells us nothing about the topology of the ambient manifold. Perhaps, however, if we first repeatedly reversed our two constructions, popping all tetrahedral and prismatic bubbles in a foam, the value of z for the resulting simpler triangulation might give some topological information. Presumably also, knowledge that every edge has valence no more than 5 (stronger than the assumption $n \leq 5$) is enough to conclude that the manifold is spherical.

Aste, Boosé and Rivier [3] define a quantity n^* , similar to n , and relate this to z with an intriguing argument which looks at the growth of the number of bubbles in layers outward from a central bubble in the foam. In terms of the triangulation, we can outline their idea as follows. Number the vertices such that all neighbors of a vertex labeled $t - 1$ are labeled at most t . For instance, we might start with one vertex labeled 0, then label its neighbors 1, and so on. Any edge, triangle or tetrahedron now has its vertices all labeled t or $t + 1$ for some t . We will say it has type t^k if exactly k of the vertices are labeled t .

The main assumption of [3] now is that there are no tetrahedra of type t^4 . This should mean that the t^3 triangles fit together into a surface L_t , called a layer of the triangulation. Edges of type t^1 go from this layer to the next, and we let n^* be the average valence of such edges. Another assumption is that this average is independent of t , or at least quickly approaches a limit value n^* as t increases. Finally, we assume also that the topology of each layer is bounded (so that we don't see increasing numbers of handles as $t \rightarrow \infty$). Then the argument of [3] sets up a recursion for the number of edges and vertices within L_t . Assuming that these numbers grow quadratically with t (rather than growing exponentially or staying bounded) in a triangulation of Euclidean space, the result is that

$$z = 6 + \frac{8}{n^* - 4}.$$

As we increase z , this equation shows that n^* decreases, while (*) shows that n increases. Naively, we might expect $n = n^*$ in most foams, since it is hard to see why edges between layers should be different from those within layers. By the monotonicity of n and n^* means that there is only one solution to $n = n^*$, namely

$$n = (10 + 2\sqrt{7})/3 \approx 5.097, \quad z = 8 + 2\sqrt{7} \approx 13.29;$$

note that these values are less than n_0 and z_0 . Actually, most known foams (including all the TCP foams of Section 7) have values greater than these, suggesting that they have $n^* < n$.

For the BCC triangulation dual to Kelvin's foam (which has $z = 14$ and $n = 5\frac{1}{7}$), if we start by giving label 0 to all vertices in the $z = 0$ plane, then we find label 1 is assigned both to the shifted vertices at $z = 1$ and to the vertices directly above at $z = 2$. Moving upwards, each label $t > 0$ is assigned to such a "double layer". This means that there are many t^4 tetrahedra, so [3] does not apply, and indeed $n^* = 5\frac{1}{3}$ does not satisfy their equation.

But if we instead start with label 0 on all vertices in the diagonal plane $x + y + z = 0$, then parallel planes have increasing t , and everything works out. For this layering, $n^* = 5$, which agrees with $z = 14$. Note also that if we start with a single cell numbered 0, then as we move outward, the layers look like octahedra. That is, they grow to essentially behave like the diagonal layers with $n^* = 5$. We might hope that other periodic foams will have the same property. That is, if we start with a single cell, and build successive layers around it, perhaps they will naturally find the appropriate directions, to give a value of n^* fitting with the prediction of [3].

In this context, we would like to suggest a reason why it might be natural to expect $n^* < n$. In a slightly distorted regular tetrahedron, we expect that longer edges will tend to have greater dihedral angles, and hence a smaller valences. However, longer edges are also more likely to reach into a new layer, that is, to enter in the calculation of n^* . If they have smaller than average valence, then indeed $n^* < n$, as observed for the Kelvin foam and predicted for any foam with $z > 13.3$.

7. TCP structures

As we have observed, the tetrahedral angle $\arccos(-1/3) \approx 109.47^\circ$, found at foam corners, is between the average angle of a pentagon (108°) and that of a hexagon (120°). We are led to consider foams with only pentagonal or hexagonal faces.

Chemists have studied transition metal alloys in which the atoms pack in nearly regular tetrahedra. These tetrahedrally close-packed (TCP) structures were first described by Frank and Kasper [11, 12], and have been studied extensively by the Shoemakers [46] among others. In all of these structures, the Voronoi cell of each atom has one of four combinatorial types; these are exactly the four polyhedra which have only pentagonal and hexagonal faces, with no adjacent hexagons. Mathematically, we can define of TCP structures in this way [26]: triangulations whose combinatorial duals (called TCP foams) have only these four types of cells. Since this

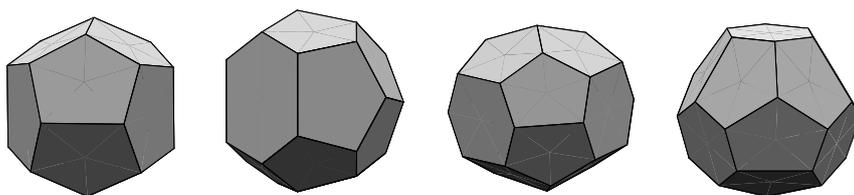


Figure 6. There are four types of cells found in TCP foams, each with 12 pentagonal faces like the dodecahedron at the left. The remaining three have in addition two, three or four hexagonal faces, arranged antipodally, equatorially, and tetrahedrally.

definition deals only with the combinatorics of the triangulation, presumably it allows some examples with too much geometric distortion to work chemically as TCP structures.

Each of the four types of cells found in TCP foams has 12 pentagonal faces, as seen in Fig. 6. One type is the (pentagonal) dodecahedron. Cells of the other three types have additionally two, three, or four hexagonal faces, which are arranged antipodally, equatorially, or tetrahedrally (respectively). The 14-hedron can be viewed as a “ $\frac{6}{5}$ -fold unwrapping” of a dodecahedron, since it has six-fold symmetry through the centers of the opposite hexagons.

Perhaps some geometric condition on the tetrahedra (saying that their edge lengths or dihedral angles are close to equal) would guarantee that the dual cells could only be one of the four TCP types. Note however, that edge lengths differ less in BCC tetrahedra (where they are 2 and $\sqrt{3}$) than in the TCP structure A15 (where they are 2, $\sqrt{5}$ and $\sqrt{6}$).

It is an interesting open question just what foams are possible with these four types of cells. Dually, we are asking for triangulations of space in which each edge has valence five or six, and no triangle has two edges of valence six. We know we cannot use only dodecahedra, since a foam made of these alone (meeting tetrahedrally) will fill a spherical space, not \mathbb{R}^3 .

This seems to be the only mathematical theorem restricting how these cells can fit together, but the chemists have made interesting additional observations. We can define three *basic* periodic TCP structures. A15, observed for instance in Cr_3Si , has one 12-hedron and three 14-hedra in a fundamental domain. Z, observed in Zr_4Al_3 , has three 12-hedra, two 14-hedra, and two 15-hedra. And C15, observed by Friauf and Laves in MgCu_2 , has two 12-hedra and one 16-hedron. (These are shown in Fig. 7 and described in more detail in Section 8 below.)

If we describe potential TCP structures by the ratio of cells they have of the four types, we could plot them all within an abstract tetrahedron, as convex combinations of the four vertices. But the observation of Yarmolyuk and Kripyakevich [57] is that all known TCP structures are, in fact, convex combinations of the three basic ones just mentioned, so they all get plotted

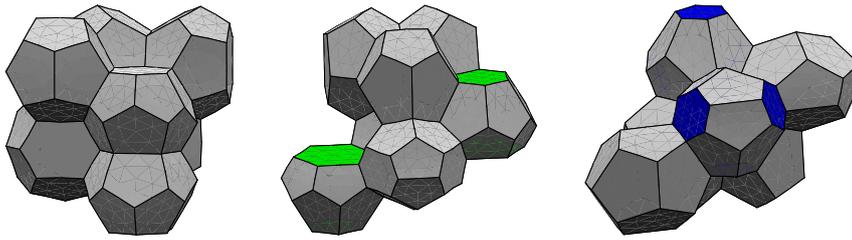


Figure 7. The three basic TCP structures. A15, left, has two 12-hedra (one upper-left) and six 14-hedra (in columns like lower right) in a cubic cell. Z, center, has three 12s (top), two 14s (in vertical columns but separated in the figure, with unequal horizontal hexagons), and three 15s (with vertical hexagons, one lower front). C15, right, has four 12s and four 16s (with hexagons darkened) in a cubic cell.

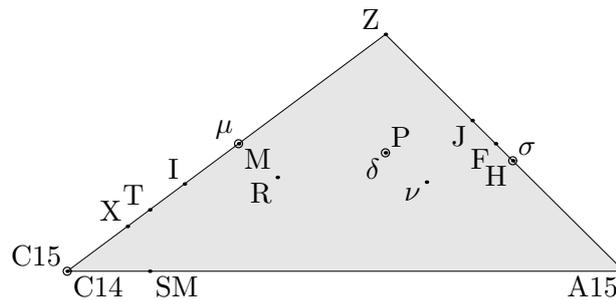


Figure 8. The basic TCP structures A15, Z, and C15 lie at the corners of this triangle, where the horizontal axis plots the average number of faces per cell (ranging from $13\frac{1}{2}$ to $13\frac{3}{2}$) and the vertical axis plots the fraction of cells which are 15-hedra (ranging from 0 to $\frac{2}{7}$). Every known TCP structure, when described by its numbers of 12-, 14-, 15-, and 16-sided cells, is a convex combination of these three. Distinct structures may appear at the same point, but need not then have the same properties.

within the triangle of Fig. 8. As we mentioned before, the Shoemakers [46] attempt to explain this observation by noting that even for a tiling by tetrahedra somewhat distorted from regular, we expect to have $z \approx z_0$. Certainly the three basic TCP structures, with $z = 13\frac{1}{2}$ for A15, $z = 13\frac{3}{7}$ for Z, and $z = 13\frac{1}{3}$ for C15, are close to this value. We will describe in the next section how to construct some convex combinations of these basic structures; the resulting foams will also have $z \approx z_0$.

It is interesting to note that there are other chemical structures, which exhibit the structure of the TCP Voronoi cells (that is, the dual TCP foams) more explicitly. These are described in detail by O’Keeffe [36], but we outline the ideas here. In clathrates, large gas molecules are trapped (at the location of the TCP sites) inside water cages: oxygen atoms sit at the Voronoi corners, bonded by hydrogen along the borders (Voronoi edges). For instance, chlorine hydrate $(\text{Cl}_2)_6 \cdot 46\text{H}_2\text{O}$, has the structure of Voronoi cells for A15, called a Type I clathrate. Many salt hydrates have instead the struc-

ture of the Z foam (called Type III), while others use non-TCP structures like the BCC cells of the Kelvin foam. Some zeolites have similar structures, with Si (and possibly some Al) atoms at the corners, bonded tetrahedrally to oxygen atoms along the borders. For instance sodalite $\text{Na}_4\text{Al}_3\text{Si}_3\text{O}_{12}\text{Cl}$ has the BCC foam structure, with Cl atoms at the Voronoi sites, and Na in the centers of the hexagonal faces. Sodium silicide, when heated properly, can generate silicon cages in the A15 or C15 foam pattern (Type I or Type II clathrates) with sodium trapped inside. It should be noted, however, that in other zeolites the 4-connected silicate nets are less dense and do not close up into cages. (See [37] for more information on many interesting crystal structures, and [29] for instance for the structures of zeolites.)

8. Some constructions for TCP foams

Perhaps the most common partition of space into equal-volume cells is by cubes, which form the Voronoi cells for the simple cubic lattice. Of course this is a degenerate case of the Voronoi construction, and the cells fail to meet tetrahedrally, so this does not make a good foam. These degenerate cubic corners (at the “holes” in the lattice) themselves form another cubic lattice. The two taken together form the BCC lattice. This structure, of course is the one whose Voronoi cells give rise to the Kelvin foam.

We might consider now repeating the process, adding some holes as new sites. The BCC lattice has only one kind of Voronoi corner (or hole). If the lattice is scaled so neighboring sites are at distance 4 and $2\sqrt{3}$ then the hole is at a distance of $\sqrt{5}$ from the nearest sites. Adjacent corners are much closer to each other, at distance $\sqrt{2}$, so if we added them all as sites, we would not be close to having regular Delone tetrahedra. But because all the faces in the foam have even numbers of sides (4 or 6), these corners can be colored black and white so that adjacent corners are different colors. If we look now just at black corners, adjacent ones are at distance 2 or $\sqrt{6}$. The structure with Si atoms at the BCC lattice points, and Cr atoms at these black corners, is the A15 structure of Cr_3Si . Its Voronoi decomposition has cells of two types (12- and 14-hedra) and forms the TCP foam of Weaire and Phelan [53], corresponding to a Type I clathrate.

If we start with the face-centered cubic lattice, it has rhombic dodecahedra for its Voronoi cells. The corners include the deep holes and (up to lattice translation) two kinds of shallow holes. The original lattice together with one kind of shallow hole forms the diamond network. If we take all its points as sites and repeat the Voronoi construction, the new corners include all the other original holes, plus certain new holes. The C15 structure for MgCu_2 has Mg atoms in the diamond network and Cu atoms at these new holes. Its Voronoi diagram in turn is the C15 foam, corresponding to

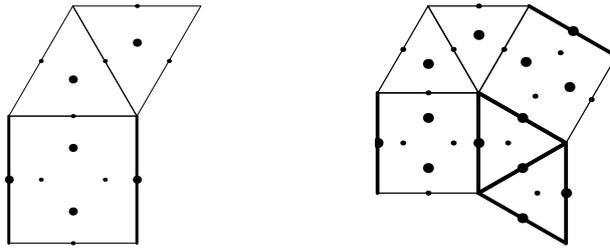


Figure 9. Here are two examples (σ , left, and H , right) of the construction for TCP structures from tilings by squares and triangles. In each case we see only a fundamental domain. The thick and thin lines correspond to the red and blue edges in the text. Small dots are sites at height $4k + 1$, large dots are sites at height $4k - 1$, and the square and triangle vertices have sites at height $2k$.

a Type II clathrate. We could construct Z in a similar fashion, filling in certain holes in a lattice of hexagonal prisms.

The work of the Frank and Kasper suggests the following construction for infinitely many new TCP structures, as convex combinations of A15 and Z . Consider an arbitrary tiling of the plane by squares and equilateral triangles. (The regular square 4^2 and triangular 3^6 tilings will give A15 and Z , respectively.) We will describe the location of sites in space for a corresponding TCP structure. Over each vertex we will find a vertical stack of 14-hedral cells, sharing hexagons, and centered at even integer heights.

Each edge of the tiling makes an angle some multiple of 30° with the horizontal. We mark it red or blue depending on whether this is an odd or even multiple, as in Fig. 9. The edge colors are the same around any triangle, but alternate around a square. At odd heights we find two different kinds of layers, with centers of 12-, 14-, and 15-hedral cells. The layers at height $4k + 1$ have sites at the midpoints of the blue edges, at the centers of the red triangles, and also halfway between the center of any square and each of its red edges. The layers at height $4k - 1$ are constructed in the same fashion, after interchanging red and blue.

There are infinitely many possibilities here; if the tiling we start with has triangles and squares in the ratio $2a : b$, then the foam obtained has 12-, 14- and 15-hedral cells in the ratio $3a + 2b : 2a + 6b : 2a$, which is just the appropriate convex combination of the ratios for Z and A15. Periodic tilings lead to periodic structures; the simplest ones have been observed in nature as TCP structures. For instance, the semi-regular tiling with vertices 3^34^2 (which alternates layers of squares and triangles) gives the H structure, while the one with vertices 3^2434 (the snub square tiling) gives the σ structure.

Sadoc and Mosseri [43] have suggested a refinement procedure for a three-dimensional triangulation which does not distort the shapes of the

tetrahedra very much. Applying this procedure to a TCP structure leads to another one. For instance, starting from the A15 structure, they derived the TCP structure called SM, which at the time was unknown in nature, although it is quite similar to the T structure previously observed.

Their construction adds two new vertices along each edge of the triangulation (trisecting the edges), and one new vertex at the center of each tetrahedron. New edges are drawn from this center vertex to the trisecting vertices and also to the centers of adjacent tetrahedra.

Each tetrahedron is divided into four small tetrahedra (at its corners) and a truncated tetrahedron, which is coned from its center to form four more tetrahedra and four hexagonal pyramids. When the hexagonal pyramids from two adjacent tetrahedra are combined into a bipyramid, the edge connecting the two center vertices divides it into six tetrahedra.

Compared to other ways to refine three-dimensional triangulations, this construction preserves the shape of tetrahedra quite well. If we start with a triangulation having two adjacent regular tetrahedra (with side 3), then some of the tetrahedra created in this construction are regular (with side 1), but others are somewhat irregular, have some edges of length $\sqrt{11/8}$ or $\sqrt{3/2}$. Of course, adjusting the positions of the new vertices might lead to slightly less variation in new edge lengths. But no matter how nice a triangulation we start with, if we apply this construction repeatedly, it seems we may be forced to get tetrahedra further and further from regular. Thus the suggestion of [43], that such repeated refinement applied to a regular triangulation of S^3 might give a Euclidean TCP structure, seems questionable.

9. Kelvin's problem

Over one hundred years ago, Kelvin proposed the problem of partitioning space into equal-volume cells using the least interface area per cell [50]. He suggested that the solution might be what we have been calling the Kelvin foam, a relaxation of the Voronoi diagram for the BCC lattice. Although Weaire and Phelan now have a better partition, Kelvin's is still conjectured [48] to be the best if the cells are required to be congruent or to have equal pressure. Mathematically, there is no theory to suggest that such a best infinite cluster should exist, but we expect that it will, and will have the structure of a foam.

Weaire and Phelan [53] were the first to consider using TCP structures as foams. Although they initially thought of these as models for wet foams, they quickly discovered that the (dry) A15 foam was a more efficient partition of space than Kelvin's candidate. To go beyond their good numerical evidence (from Brakke's Evolver) and give a rigorous proof that their foam

is better than Kelvin's, we need to give a bound on how much the Kelvin foam can relax.

Although, as we have mentioned, there is no general mathematical theory for the relaxation step in constructing foams from Voronoi cells, for Kelvin's foam there is enough symmetry that we need only consider mean-curvature flow on a single surface. In Fig. 1, right, we saw a symmetric unit of Kelvin's foam, bounded by lines of rotational symmetry. In that picture, the vertical sheet is in a mirror plane, so we need only solve for one of the other sheets, a minimal surface with two fixed boundary lines and one free boundary with 120° contact angle. We can use this analysis to show that a unique foam exists in Kelvin's pattern, and a slicing argument then gives a lower bound on its area [26]. This bound, it turns out, suffices to prove that even the unrelaxed Weaire-Phelan A15 foam beats Kelvin.

Given this example, it is natural to look for good equal-volume partitions among the other TCP foams. Rivier [41] proposed that since C15 has (among the known TCP structures) the lowest z and thus the highest proportion of pentagons, it might give an even better partition. Instead, computer experiments [22] suggest that among all TCP foams, A15 is the most efficient, and C15 among the least. (Here, we are looking at equal-volume foams. For each pattern (like Z) without cubic symmetry, we adjust the lattice parameters to get rid of any stress tensor; this ensures that we have the most efficient partition in that pattern.) Probably the 12-hedral and 16-hedral cells of C15 naturally have such different sizes that distorting them to make the volumes equal ruins whatever advantage pentagons give. (If we considered a modified Kelvin problem where one-third of the cells were to have somewhat larger volume, then presumably the C15 structure would do very well.) The author has also made computer experiments with equal-volume foams generated from other chemical structures; the foam from γ -brass, for instance, is better than that from C15, despite having even some triangular faces. Perhaps pentagons are not as desirable in foams as has been assumed.

References

1. Alexandrov, A. D.: 1958, 'Uniqueness Theorems for Surfaces in the Large, I'. *Vestnik Leningrad Univ. Math.* **19**(13), 5–8. English transl. in Amer. Math. Soc. Transl. (Ser. 2) **21** (1962), 412–416.
2. Almgren, Jr., F. J.: 1976, 'Existence and Regularity Almost Everywhere of Solutions to Elliptic Variational Problems with Constraints'. *Mem. Amer. Math. Soc.* **4**(165).
3. Aste, T., D. Boosé, and N. Rivier: 1996, 'From One Cell to the Whole Froth: A Dynamical Map'. *Phys. Rev.* **53**, 6181–6191.
4. Brakke, K. and F. Morgan: 1996, 'Instability of the Wet X Soap Film'. Preprint.
5. Brakke, K. A.: 1992, 'The Surface Evolver'. *Exper. Math.* **1**(2), 141–165.
6. Brakke, K. A. and J. M. Sullivan: 1997, 'Using Symmetry Features of the Surface Evolver to Study Foams'. In: K. Polthier and H.-C. Hege (eds.): *Visualization and*

- Mathematics*. Heidelberg, pp. 95–117.
7. Choe, J.: 1989, ‘On the Existence and Regularity of Fundamental Domains with Least Boundary Area’. *J. Diff. Geom.* **29**, 623–663.
 8. Coxeter, H. S. M.: 1958, ‘Close-Packing and Froth’. *Ill. J. Math.* **2**(4B), 746–758. Reprinted in [52].
 9. Delaunay, C.: 1841, ‘Sur la surface de révolution, dont la courbure moyenne est constante’. *Journal de mathématiques* **6**, 309–320.
 10. Dierkes, U., S. Hildebrandt, A. Küster, and O. Wohlrab: 1992, *Minimal Surfaces I*, Vol. 295 of *Grundlehren der Mathematischen Wissenschaften*. Berlin: Springer-Verlag.
 11. Frank, F. C. and J. S. Kasper: 1958, ‘Complex Alloy Structures Regarded as Sphere Packings. I. Definitions and Basic Principles’. *Acta Crystall.* **11**, 184–190.
 12. Frank, F. C. and J. S. Kasper: 1959, ‘Complex Alloy Structures Regarded as Sphere Packings. II. Analysis and Classification of Representative Structures’. *Acta Crystall.* **12**, 483–499.
 13. Große-Brauckmann, K., R. Kusner, and J. M. Sullivan: 1997, ‘Classification of Embedded Constant Mean Curvature Surfaces with Genus Zero and Three Ends’. GANG Preprint IV.29, UMass.
 14. Große-Brauckmann, K., R. Kusner, and J. M. Sullivan: 1998, ‘Constant Mean Curvature Surfaces with Cylindrical Ends’. To appear in the Springer proceedings of VisMath’97.
 15. Hass, J., M. Hutchings, and R. Schlafly: 1995, ‘The Double Bubble Conjecture’. *Electron. Res. Announc. Amer. Math. Soc.* **1**(3), 98–102.
 16. Hildebrandt, S.: 1970, ‘On the Plateau problem for surfaces of constant mean curvature’. *Comm. Pure Appl. Math.* **23**, 97–114.
 17. Hildebrandt, S. and A. Tromba: 1996, *The Parsimonious Universe*. New York: Copernicus.
 18. Hoffman, D. and W. H. Meeks, III: 1990, ‘Embedded Minimal Surfaces of Finite Topology’. *Ann. of Math.* **131**(1), 1–34.
 19. Hsu, L., R. Kusner, and J. M. Sullivan: 1992, ‘Minimizing the Squared Mean Curvature Integral for Surfaces in Space Forms’. *Experimental Mathematics* **1**(3), 191–207.
 20. Jülicher, F., U. Seifert, and R. Lipowsky: 1993, ‘Conformal Degeneracy and Conformal Diffusion of Vesicles’. *Phys. Rev. Lett.* **71**, 452–455.
 21. Korevaar, N., R. Kusner, and B. Solomon: 1989, ‘The Structure of Complete Embedded Surfaces with Constant Mean Curvature’. *J. Diff. Geom.* **30**, 465–503.
 22. Kraynik, A. M., R. Kusner, R. Phelan, and J. M. Sullivan, ‘TCP Structures as Equal-Volume Foams’. In preparation.
 23. Kraynik, A. M. and D. A. Reinelt: 1996, ‘Elastic-Plastic Behavior of a Kelvin Foam’. *Forma* **11**(3), 255–270. Reprinted in [52].
 24. Kusner, R.: 1992, ‘The Number of Faces in a Minimal Foam’. *Proc. R. Soc. Lond.* **439**, 683–686.
 25. Kusner, R. and N. Schmitt: 1996, ‘On the Spinor Representation of Minimal Surfaces’. GANG preprint III.27, UMass.
 26. Kusner, R. and J. M. Sullivan: 1996, ‘Comparing the Weaire-Phelan Equal-Volume Foam to Kelvin’s Foam’. *Forma* **11**(3), 233–242. Reprinted in [52].
 27. Luo, F. and R. Stong: 1993, ‘Combinatorics of Triangulations of 3-Manifolds’. *Trans. Amer. Math. Soc.* **337**(2), 891–906.
 28. Matzke, E. B.: 1946, ‘The Three-Dimensional Shape of Bubbles in Foam’. *Amer. J. Botany* **33**, 58–80.
 29. Meier, W. M. and D. H. Olson: 1992, *Atlas of Zeolite Structure Types*. Butterworths, 3rd edition.
 30. Michalet, X. and D. Bensimon: 1995, ‘Observations of Stable Shapes and Conformal Diffusion in Genus 2 Vesicles’. *Science* **269**, 666–668.
 31. Morgan, F.: 1994, ‘Clusters Minimizing Area Plus Length of Singular Curves’. *Math. Ann.* **299**(4), 697–714.

32. Morgan, F.: 1995a, *A Beginner's Guide to Geometric Measure Theory*. Academic Press, 2nd edition.
33. Morgan, F.: 1995b, 'The Double Soap Bubble Conjecture'. *MAA FOCUS* pp. 6–7. Dec. 1995.
34. Morgan, F.: 1996, 'The Hexagonal Honeycomb Conjecture'. Preprint.
35. Okabe, A., B. Boots, and K. Sugihara: 1992, *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. Wiley & Sons.
36. O'Keeffe, M.: 1997, 'Crystal Structures as Periodic Foams and *vice versa*'. Appearing in this volume.
37. O'Keeffe, M. and B. G. Hyde: 1996, *Crystal Structures I: Patterns and Symmetry*. Washington: Mineral Soc. Amer.
38. Osserman, R.: 1986, *A Survey of Minimal Surfaces*. New York: Dover Publications, 2nd edition.
39. Plateau, J.: 1873, *Statique Expérimentale et Théorique des Liquides Soumis aux Seules Forces Moléculaires*. Paris: Gauthier-villars.
40. Reinelt, D. A. and A. M. Kraynik: 1993, 'Large Elastic Deformations of Three-Dimensional Foams and Highly Concentrated Emulsions'. *J. of Colloid and Interface Science* **159**, 460–470.
41. Rivier, N.: 1994, 'Kelvin's Conjecture on Minimal Froths and the Counter-Example of Weaire and Phelan'. *Europhys. Lett.* **7**(6), 523–528.
42. Rogers, C. A.: 1958, 'The Packing of Equal Spheres'. *Proc. London Math. Soc.* **8**, 609–620.
43. Sadoc, J.-F. and R. Mosseri: 1982, 'Order and Disorder in Amorphous Tetrahedrally Coordinated Semiconductors: A Curved-Space Description'. *Philos. Mag.* **45**, 467.
44. Schwarz, H. A.: 1884, 'Beweis des Satzes, dass die Kugel kleinere Oberfläche besitzt, als jeder andere Körper gleichen Volumnes'. *Nach. Ges. Wiss. Göttingen* pp. 1–13. Reprinted in 1972 in *Gesammelte mathematische Abhandlungen*, pp. II.327–340, New York: Chelsea.
45. Senechal, M.: 1990, *Crystalline Symmetries*. Adam Hilger.
46. Shoemaker, D. P. and C. B. Shoemaker: 1986, 'Concerning the Relative Numbers of Atomic Coordination Types in Tetrahedrally Close Packed Metal Structures'. *Acta Crystall.* **42**, 3–11.
47. Sullivan, J. M.: 1988, 'The *vcs* Software for Computing Voronoi Diagrams'. Available by email from jms@math.uiuc.edu.
48. Sullivan, J. M. and F. Morgan (Editors): 1996, 'Open Problems in Soap Bubble Geometry: Posed at the Burlington Mathfest in August 1995'. *International J. of Math.* **7**(6), 833–842.
49. Taylor, J. E.: 1976, 'The Structure of Singularities in Soap-Bubble-Like and Soap-Film-Like Minimal Surfaces'. *Ann. of Math.* **103**, 489–539.
50. Thompson, Sir W. (Lord Kelvin): 1887, 'On the Division of Space with Minimum Partitional Area'. *Philos. Mag.* **24**, 503–514. Also published in *Acta Math.* **11**, 121–134, and reprinted in [52].
51. Thurston, W. P.: 1997, *Three-Dimensional Geometry and Topology*, Vol. 1. Princeton. Edited by Silvio Levy.
52. Weaire, D. (ed.): 1997, *The Kelvin Problem*. Taylor & Francis.
53. Weaire, D. and R. Phelan: 1994, 'A Counter-Example to Kelvin's Conjecture on Minimal Surfaces'. *Phil. Mag. Lett.* **69**(2), 107–110. Reprinted in [52].
54. Williams, R. E.: 1968, 'Space Filling Polyhedron: Its Relation to Aggregates of Soap Bubbles, Plant Cells, and Metal Crystallites'. *Science* **161**, 276–277.
55. Willmore, T. J.: 1992, 'A Survey on Willmore Immersions'. In: *Geometry and Topology of Submanifolds, IV (Leuven, 1991)*. pp. 11–16.
56. Wintz, W., H.-G. Döbereiner, and U. Seifert: 1996, 'Starfish Vesicles'. *Europhys. Lett.* **33**, 403–408.
57. Yarmolyuk, Y. P. and P. I. Kripyakevich: 1974. *Kristallografiya* **19**, 539–545. Translated in *Sov. Phys. Crystallogr.* **19**, 334–337.