

Fat 4-polytopes and fatter 3-spheres

David Eppstein,^{1,*} Greg Kuperberg,^{2,†} and Günter M. Ziegler^{3,‡}

¹Department of Information and Computer Science, University of California, Irvine, CA 92697

²Department of Mathematics, University of California, Davis, CA 95616

³Institut für Mathematik, MA 6-2, Technische Universität Berlin, D-10623 Berlin, Germany

We introduce the *fatness* parameter of a 4-dimensional polytope P , defined as $\phi(P) = (f_1 + f_2)/(f_0 + f_3)$. It arises in an important open problem in 4-dimensional combinatorial geometry: Is the fatness of convex 4-polytopes bounded?

We describe and analyze a hyperbolic geometry construction that produces 4-polytopes with fatness $\phi(P) > 5.048$, as well as the first infinite family of 2-simple, 2-simplicial 4-polytopes. Moreover, using a construction via finite covering spaces of surfaces, we show that fatness is not bounded for the more general class of strongly regular CW decompositions of the 3-sphere.

1. INTRODUCTION

The characterization of the set \mathcal{F}_3 of f -vectors of convex 3-dimensional polytopes (from 1906, due to Steinitz [28]) is well-known and explicit, with a simple proof: An integer vector (f_0, f_1, f_2) is the f -vector of a 3-polytope if and only if it satisfies

- $f_1 = f_0 + f_2 - 2$ (the Euler equation),
- $f_2 \leq 2f_0 - 4$ (with equality for simplicial polytopes), and
- $f_0 \leq 2f_2 - 4$ (with equality for simple polytopes).

(Recall that by the definition of the f -vector, f_k is the number of k -faces of the polytope.) This simple result is interesting for several reasons:

- The set of f -vectors is the set of *all* the integer points in a closed 2-dimensional *polyhedral cone* (whose apex is the f -vector $f(\Delta_3) = (4, 6, 4)$ of a 3-dimensional simplex). In particular, it is *convex* in the sense that $\mathcal{F}_3 = \text{conv}(\mathcal{F}_3) \cap \mathbb{Z}^3$.
- The same characterization holds for convex 3-polytopes (geometric objects), more generally for strongly regular CW 2-spheres (topological objects), and yet more generally for Eulerian lattices of length 4 (combinatorial objects [27]).

In contrast to this explicit and complete description of \mathcal{F}_3 , our knowledge of the set \mathcal{F}_4 of f -vectors of (convex) 4-polytopes (see Bayer [2] and Höppner and Ziegler [14]) is very incomplete. We know that the set \mathcal{F}_4 of all f -vectors of 4-dimensional polytopes has no similarly simple description. In particular, the convex hull of \mathcal{F}_4 is not a cone, it is not a closed set, and not all integer points in the convex hull are f -vectors. Also, the 3-dimensional cone with apex $f(\Delta_4)$ spanned by \mathcal{F}_4 is not closed, and its closure may not be polyhedral.

Only the two extreme cases of simplicial and of simple 4-polytopes (or 3-spheres) are well-understood. Their f -vectors correspond to faces of the convex hull of \mathcal{F}_4 , defined by the valid inequalities $f_2 \geq 2f_3$ and $f_1 \geq 2f_0$, and the g -Theorem, proved for 4-polytopes by Barnette [1] and for 3-spheres by Walkup [34], provides *complete* characterizations of their f -vectors. (The g -Theorem for general simplicial polytopes was famously conjectured by McMullen [21] and proved by Billera and Lee [4] and Stanley [26]. See [35, §8.6] for a review.)

But we have no similarly complete picture of other extremal types of 4-polytopes. In particular, we cannot currently answer the following key question: Is there a constant c such that all 4-dimensional convex polytopes P satisfy the inequality

$$f_1(P) + f_2(P) \leq c(f_0(P) + f_3(P))?$$

To study this question, we introduce the *fatness* parameter

$$\phi(P) \stackrel{\text{def}}{=} \frac{f_1(P) + f_2(P)}{f_0(P) + f_3(P)}$$

of a 4-polytope P . We would like to know whether fatness is bounded.

For example, the 4-simplex has fatness 2, while the 4-cube and the 4-cross polytope have fatness $\frac{56}{24} = \frac{7}{3}$. More generally, if P is simple, then we can substitute the Dehn-Sommerville relations

$$f_2(P) = f_1(P) + f_3(P) - f_0(P) \quad f_1(P) = 2f_0(P)$$

into the formula for fatness, yielding

$$\phi(P) = \frac{f_1(P) + f_2(P)}{f_0(P) + f_3(P)} = \frac{3f_0(P) + f_3(P)}{f_0(P) + f_3(P)} < 3.$$

Since every 4-polytope and its dual have the same fatness, the same upper bound holds for simplicial 4-polytopes. On the other hand, the “neighborly cubical” 4-polytopes of Joswig and Ziegler [16] have f -vectors

$$(4, 2n, 3n - 6, n - 2) \cdot 2^{n-2},$$

and thus fatness

$$\phi = \frac{5n - 6}{n + 2} \rightarrow 5.$$

*Email: eppstein@ics.uci.edu; Supported in part by NSF grant CCR #9912338

†Email: greg@math.ucdavis.edu; Supported by NSF grant DMS #0072342

‡Email: ziegler@math.tu-berlin.de; Supported by a DFG Leibniz grant

In particular, the construction of these polytopes disproved the conjectured flag-vector inequalities of Bayer [2, pp. 145, 149] and Billera and Ehrenborg [14, p. 109].

The main results of this paper are two lower bounds on fatness:

Theorem 1. *There are convex 4-polytopes P with fatness $\phi(P) > 5.048$.*

Theorem 2. *The fatness of cellulated 3-spheres is not bounded. A 3-sphere S with N vertices may have fatness as high as $\phi(S) = \Omega(N^{1/12})$.*

We will prove Theorem 1 in Section 3 and Theorem 2 in Section 4, and present a number of related results along the way.

2. CONVENTIONS

Let X be a finite CW complex. If X is identified with a manifold M , it is also called a *cellulation* of M . The complex X is *regular* if all closed cells are embedded [22, §38]. If X is regular, we define it to be *strongly regular* if in addition the intersection of any two closed cells is a cell. For example, every simplicial complex is a strongly regular CW complex. The complex X is *perfect* if the boundary maps of its chain complex vanish. (A non-zero-dimensional perfect complex is never regular.)

The *f-vector* of a cellulation X , denoted $f(X) = (f_0, f_1, \dots)$, counts the number of cells in each dimension: $f_0(X)$ is the number of vertices, $f_1(X)$ is the number of edges, etc.

If X is 2-dimensional, we define its fatness as

$$\phi(X) \stackrel{\text{def}}{=} \frac{f_1(X)}{f_0(X) + f_2(X)}.$$

If X is 3-dimensional, we define its fatness X as

$$\phi(X) \stackrel{\text{def}}{=} \frac{f_1(X) + f_2(X)}{f_0(X) + f_3(X)}.$$

If P is a convex d -polytope, then its *f-vector* $f(P)$ is defined to be the *f-vector* of its boundary complex, which is a strongly regular $(d-1)$ -sphere. If $d=4$ we can thus consider $\phi(P)$, the fatness of P . The faces of P of dimension 0 and 1 are called *vertices* and *edges* while the faces of dimension $d-1$ and $d-2$ are called *facets* and *ridges*. We extend this terminology to general cellulations of $(d-1)$ -manifolds.

The *flag vector* of a regular cellulation X , and likewise the flag vector of a polytope P , counts the number of nested sequences of cells with prespecified dimensions. For example, $f_{013}(X)$ is the number triples consisting of a 3-cell of X , an edge of the 3-cell, and a vertex of the edge; if P is a 4-cube, then

$$f_{013}(P) = 192.$$

A convex polytope P is *simplicial* if each facet of P is a simplex. It is *simple* if its polar dual P^Δ is simplicial, or equivalently if the cone of each vertex matches that of a simplex.

3. 4-POLYTOPES

In this section we construct families of 4-polytopes with several interesting properties:

- They are the first known infinite families of *2-simple*, *2-simplicial* 4-polytopes, that is, polytopes in which all 2-faces and all dual 2-faces are triangles (all edges are “co-simple”). *E.g.*, Bayer [2] says that it would be interesting to have an infinite family.

As far as we know, there were only six such polytopes previously known: the simplex, the hypersimplex (the set of points in $[0, 1]^4$ with coordinate sum between 1 and 2), the dual of the hypersimplex, the 24-cell, and a gluing of two hypersimplices (due to Braden [6]), and the dual of the gluing. (There are claims in Grünbaum [13, p. 82, resp. p. 170] that results of Shephard, resp. Perles and Shephard, imply the existence of infinitely many 2-simple 2-simplicial 4-polytopes. Both claims appear to be incorrect.)

- They are the fattest known convex 4-polytopes.
- They yield finite packings of (not necessarily congruent) spheres in \mathbb{R}^3 with slightly higher average kissing numbers than previously known examples [18].

Let $Q \subset \mathbb{R}^4$ be a 4-polytope that contains the origin in its interior. If an edge e of Q is tangent to the unit sphere $S^3 \subset \mathbb{R}^4$ at a point t , then the corresponding ridge (2-dimensional face) $F = e^\diamond$ of the polar dual $P = Q^\Delta$ is also tangent to S^3 at t . (Recall that the polar dual is defined as

$$Q^\Delta \stackrel{\text{def}}{=} \{p \mid p \cdot q \leq 1 \forall q \in Q\}.$$

Furthermore, the affine hulls of e and F form orthogonal complements in the tangent space of S^3 , so the convex hull $\text{conv}(e \cup F)$ is an orthogonal bipyramid tangent to S^3 (*cf.*, Schulte [24, Thm. 1]).

We will construct polytopes E by what we call the *E-construction*. This means that they are convex hulls

$$E \stackrel{\text{def}}{=} \text{conv}(Q \cup P),$$

where Q is a simplicial 4-dimensional polytope whose edges are tangent to the unit 3-sphere S^3 , and P is the polar dual of Q . Thus P is simple and its ridges are tangent to S^3 . (We call Q *edge-tangent* and P *ridge-tangent*.)

Proposition 3. *If P is a simple, ridge-tangent 4-polytope, then the 4-polytope $E = \text{conv}(P \cup Q)$ produced by the *E-construction* is 2-simple and 2-simplicial, with *f-vector**

$$f(E) = (f_2(P), 6f_0(P), 6f_0(P), f_2(P)),$$

and fatness

$$\phi(E) = 6 \frac{f_0(P)}{f_2(P)}.$$

Proof. Another way to view the E -construction is that E is produced from P by adding the vertices of Q sequentially. At each step, we cap a facet of P with a pyramid whose apex is a vertex of Q . Thus the new facets consist of pyramids over the ridges of P , where two pyramids with the same base (appearing in different steps) lie in the same hyperplane (tangent to S^3), and together form a bipyramid. The facets of the final polytope E are orthogonal bipyramids over the ridges of P and are tangent to S^3 . Since the 2-faces of E are pyramids over the edges of P , E is 2-simplicial.

The polytope E is 2-simple if and only if each edge is *co-simple*, i.e., contained in exactly three facets of E . The iterative construction of E shows that it has two types of edges: (i) edges of P , which are co-simple in E if and only if they are co-simple in P , and (ii) edges formed by adding pyramids, which are co-simple if and only if the facets of P are simple. Since P is simple, its facets are simple and its edges are co-simple, so E is then 2-simple.

This combinatorial description of E yields an expression for the f -vector of E in terms of the flag vector [2, 3, 14] of P . Since the facets of E are bipyramids over the ridges of P , the following identities hold:

$$\begin{aligned} f_3(E) &= f_2(P) & f_2(E) &= f_{13}(P) \\ f_1(E) &= f_1(P) + f_{03}(P) & f_0(E) &= f_0(P) + f_3(P). \end{aligned}$$

Since P is simple,

$$f_{03}(P) = 4f_0(P) \quad f_{13}(P) = 3f_1(P).$$

These identities together with Euler's equation and $f_1(P) = 2f_0(P)$ imply the proposition. \square

The f -vector of P also satisfies

$$f_0(P) - f_1(P) + f_2(P) - f_3(P) = 0 \quad f_1(P) = 2f_0(P),$$

in the second case because P is simple, so the fatness of E can also be written

$$\phi(E) = 6 \left(1 - \frac{f_3(P)}{f_2(P)} \right) = 6 \left(1 - \frac{f_0(Q)}{f_1(Q)} \right).$$

Thus maximizing the fatness of E is equivalent to maximizing the ridge-facet ratio $f_2(P)/f_3(P)$, or the average degree of the graph of Q . It also shows that the E -construction cannot achieve a fatness of 6 or more.

In light of Proposition 3, we would like to construct edge-tangent simplicial 4-polytopes. Regular simplicial 4-polytopes (suitably scaled) provide three obvious examples: the 4-simplex Δ_4 , the cross polytope C_4^Δ , and the 600-cell. From these, the E -construction produces the dual of the hypersimplex, the 24-cell, and a new 2-simple, 2-simplicial polytope with f -vector $(720, 3600, 3600, 720)$ and fatness 5, whose facets are bipyramids over pentagons.

We will construct new edge-tangent simplicial 4-polytopes by gluing together (not necessarily simplicial) edge-tangent 4-polytopes, called *atoms*, to form *compounds*. We must position the polytopes so that their facets match, they remain edge-tangent, and the resulting compound is convex. It will be very

useful to interpret the interior of the 4-dimensional unit ball as the Klein model of hyperbolic 4-space \mathbb{H}^4 , with S^3 the sphere at infinity. (See Iversen [15] and Thurston [30, Chap. 2] for introductions to hyperbolic geometry.) In particular, Euclidean lines are straight in the Klein model, Euclidean subspaces are flat, and hence any intersection of a convex polytope with \mathbb{H}^4 is a convex (hyperbolic) polyhedron. Even though the Klein model respects convexity, it does not respect angles. However, angles and convexity are preserved under hyperbolic isometries. There are enough isometries to favorably position certain 4-polytopes to produce convex compounds.

If a polytope Q is edge-tangent to S^3 , then it is hyperbolically hyperideal: Not only its vertices, but also its edges, lie beyond the sphere at infinity, except for the tangency point of each edge. Nonetheless portions of its facets and ridges lie in the finite hyperbolic realm. As a hyperbolic object the polytope Q (more precisely, $Q \cap \mathbb{H}^4$) is convex and has flat facets. The ridge r between any two adjacent facets has a well-defined hyperbolic dihedral angle, which is strictly between 0 and π if (as in our situation) the ridge properly intersects \mathbb{H}^4 . To compute this angle we can intersect r at any point t with any hyperplane R that contains the (hyperbolic) orthogonal complement to r at t . We let t be the tangency point of any edge e of the ridge, and let R be the hyperplane perpendicular to e .

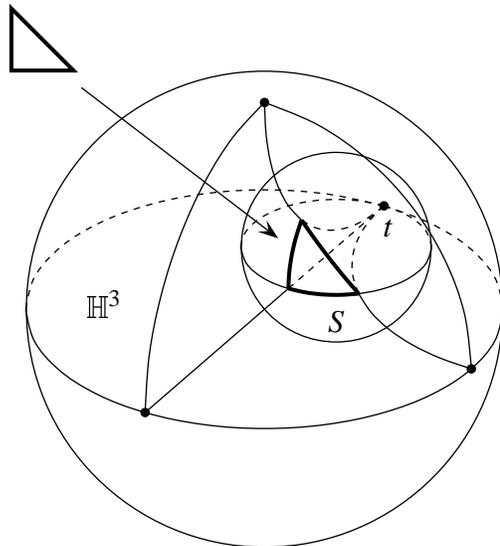


Figure 1: A cone emanating from an ideal point t in \mathbb{H}^3 in the Poincaré model, and a horosphere S incident to t . The link of t (here a right isosceles triangle) inherits Euclidean geometry from S .

Within the hyperbolic geometry of $R \cong \mathbb{H}^3$, every line emanating from the ideal point t is orthogonal to any horosphere incident to t . Thus the link of the edge e of Q is the intersection of $Q \cap R$ with a sufficiently small horosphere S at t . Since horospheres have flat Euclidean geometry [30, p. 61], the link $Q \cap S$ is a Euclidean polygon. Its edges correspond to the facets of Q that contain e , and its vertices correspond to the ridges of Q that contain e . Thus the dihedral angle of a ridge r of Q equals to the Euclidean angle of the vertex $r \cap S$ of the

Euclidean polygon $Q \cap S$. This is easier to see in the Poincaré model of hyperbolic space, because it respects angles, than in the Klein model. Figure 1 shows an example.

To summarize:

Lemma 4. *A compound of two or more polytopes is convex if and only if each ridge has hyperbolic dihedral angle less than π , or equivalently, iff each edge link is a convex Euclidean polygon.*

Compounds can also have interior ridges with total dihedral angle exactly 2π . But since all atoms of a compound are edge-tangent, compounds do not have any interior edges or vertices.

If Q is a regular polytope, then $Q \cap S$ is a regular polygon. The following lemma is then immediate:

Lemma 5. *If Q is a regular, edge-tangent, simplicial 4-polytope, then in the hyperbolic metric of the Klein model, its dihedral angles are $\pi/3$ (for the simplex), $\pi/2$ (for the cross polytope), and $3\pi/5$ (for the 600-cell).*

A hyperideal hyperbolic object, even if it is an edge-tangent convex polytope, can be unfavorably positioned so that it is unbounded as a Euclidean object (cf., Schulte [24, p. 508]). Fortunately there is always a bounded position as well:

Lemma 6. *Let Q be an edge-tangent, convex polytope in \mathbb{R}^d whose points of tangency with S^{d-1} do not lie in a hyperplane. Then there is a hyperbolic isometry h (extended to all of \mathbb{R}^d) such that $h(Q)$ is bounded.*

Proof. Let p lie in the interior of the convex hull of the edge tangencies and let f be any hyperbolic motion that moves p to the Euclidean origin in \mathbb{R}^d . Since the convex hull K of the edge tangencies of $h(Q)$ contains the origin, K^Δ is a bounded polytope that circumscribes S^{d-1} . Since K^Δ is facet-tangent where $h(Q)$ is edge-tangent, $h(Q) \subset K^\Delta$. \square

In the following we discuss three classes of edge-tangent simplicial convex 4-polytopes that are obtained by gluing in the Klein model: Compounds of simplices, then simplices and cross polytopes, and finally compounds from cut 600-cells. There are yet other edge-tangent compounds involving cross polytopes cut in half (i.e., pyramids over octahedra), 24-cells, and hypersimplices as atoms, but we will not discuss these here.

3.1. Compounds of simplices

In this section we classify compounds whose atoms are simplices. This includes all *stacked* polytopes, which are simplicial polytopes that decompose as a union of simplices without any interior faces other than facets. However compounds of simplices are a larger class, since they may have interior ridges.

Lemma 7. *Any edge-tangent d -simplex is hyperbolically regular.*

Proof. The proof is by induction on d , starting from the case $d = 2$, where the three tangency points define an ideal triangle in \mathbb{H}^2 . All ideal triangles are congruent [30, p. 83]. Since edge-tangent triangles are the polar duals of ideal triangles, they are all equivalent as well.

If $d > 2$, let B be a general edge-tangent d -simplex. On the one hand, there exists an edge-tangent simplex A which is regular both in hyperbolic geometry and Euclidean geometry. On the other hand, given the position of d of the vertices, there are at most two choices for the last vertex that produce an edge-tangent simplex, one on each side of the hyperplane spanned by the first d . By induction there exists an isometry that takes a face of B to a face of A and the remaining vertex to the same side. The edge-tangent constraint implies that this isometry takes the last vertex of B to the last vertex of A as well. \square

Proposition 8. *There are only three possible edge-tangent compounds of 4-simplices:*

- *the regular simplex,*
- *the bipyramid (a compound of two simplices that share a facet), and*
- *the join of a triangle and a hexagon (a compound of six simplices that share a ridge).*

Proof. Figure 2 shows all strictly convex polygons with unit-length edges tiled by unit equilateral triangles, or *triangle jewels*. Since the atoms of an edge-tangent compound of simplices are edge-tangent, they are hyperbolically regular by Lemma 7, and their edge links are equilateral triangles. Thus every edge link of a compound of simplices is a triangle jewel. Any three 4-simplices in a chain in such a compound share a ridge. In order to create an edge link matching Figure 2, they must extend to a ring of six simplices around the same ridge. Adding any further simplex to these six would create an edge link in the form of a triangle surrounded by three other triangles, which does not appear in Figure 2. \square

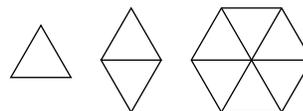


Figure 2: The 3 possible edge links of edge-tangent compounds of 4-simplices.

Two of the E -polytopes produced by Proposition 8 were previously known. If Q is the simplex, then E is dual to the hypersimplex. If Q is the bipyramid, then E is dual to Braden's glued hypersimplex. However, if Q is the six-simplex compound (dual to the product of a triangle and a hexagon), then E is a new 2-simple, 2-simplicial polytope with f -vector (27, 108, 108, 27).

Proposition 8 also implies an interesting impossibility result.

Corollary 9. *No stacked 4-polytope with more than 6 vertices is edge-tangent.*

See Schulte [24, Thm. 3] for the first examples of polytopes that have no edge-tangent realization.

3.2. Compounds of simplices and cross polytopes

Next, we consider compounds of simplices and regular cross polytopes. The edge link of any convex compound of these two types of polytopes must be one of the eleven strictly convex polygons tiled by unit triangles and squares, or *square-triangle jewels* (Figure 3). See Malkevitch [20] and Waite [33] for work on convex compounds of these shapes relaxing the requirement of strict convexity.

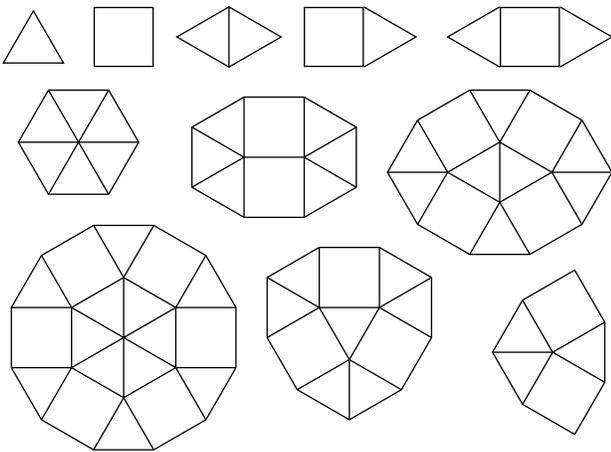


Figure 3: The 11 possible edge links of edge-tangent compounds of 4-simplices and cross polytopes.

If Q is a single cross polytope, then E is a 24-cell. We can also glue simplices onto subsets of the facets of the cross polytope. The new dihedral angles formed by such a gluing are $5\pi/6$. The resulting compound is convex as long as no two glued cross polytope facets share a ridge. We used a computer program to list the combinatorially distinct ways of choosing a subset of nonadjacent facets of the cross polytope; the results may be summarized as follows. In addition to the 24-cell, this yields 20 new 2-simple, 2-simplicial polytopes.

Proposition 10. *There are exactly 21 distinct simplicial edge-tangent compounds composed of one regular 4-dimensional cross polytope and $k \geq 0$ simplices, according to the following table:*

k	0	1	2	3	4	5	6	7	8	Total
#	1	1	3	3	6	3	2	1	1	21

We can also confirm that every square-triangle jewel arises as the edge link of an edge-tangent 4-dimensional compound.

For every jewel other than the one in the center, we can form a convex *edge bouquet* consisting of simplices and cross polytopes that meet at an edge: we replace each triangle by a simplex and each square by a cross polytope. Since the central jewel has two adjacent squares, its edge bouquet is not convex. Instead we glue two cross polytopes along a facet so that the 4 ridges of that facet are flush, *i.e.*, their dihedral angle is π . Thus we can “caulk” each such ridge with three simplices that share the ridge. The central jewel is the link of 6 of the edges of the resulting compound of 2 cross polytopes and 12 simplices.

Simplices and regular cross polytopes combine to form many other edge-tangent simplicial polytopes and hence 2-simple, 2-simplicial polytopes. In particular, these methods lead to the following theorem.

Theorem 11. *There are infinitely many combinatorially distinct 2-simple, 2-simplicial facet-tangent 4-polytopes.*

Proof. We glue n cross polytopes end-to-end. Each adjacent pair produces 4 flush ridges that we caulk with chains of three simplices. The facets to which these simplices are glued are not adjacent and so do not produce any further concavities. \square

The chain of n cross polytopes has f -vector

$$(4n + 4, 18n + 6, 28n + 4, 14n + 2).$$

Filling a concavity adds $(2, 9, 14, 7)$, so after filling the $4(n-1)$ concavities we get a simplicial polytope Q with

$$f(Q) = (12n - 4, 54n - 30, 84n - 54, 42 - 26),$$

which yields a 2-simple, 2-simplicial 4-polytope E with

$$f(E) = (54n - 30, 252n - 156, 252n - 156, 54n - 30)$$

by Proposition 3.

Remark. Every 2-simple, 2-simplicial 4-polytopes that we know is combinatorially equivalent to one which circumscribes the sphere. Are there any that are not?

We do know a few 2-simple, 2-simplicial 4-polytopes which are not E -polytopes. Trivially there is the simplex. There are a few others that arise by the fact that the 24-cell is the E -polytope of a cross polytope in 3 different ways. Color the vertices of a 24-cell red, green, and blue, so that the vertices of each color span a cross polytope. If we cap one facet of a cross-polytope by a simplex and apply the E -construction, the result is a 24-cell in which 6 facets that meet at 1 vertex are replaced by 10 facets and 4 vertices. If the replaced vertex is red, the replacement can be induced by capping either blue cross polytope or the green cross polytope and then applying the E -construction; the position of the replacement differs between the two cases. If we replace two different red vertices, one by capping the green cross polytope and the other by capping the blue cross polytope, then the resulting polytope is 2-simple and 2-simplicial but not an E -polytope. This construction has several variations: for example, we can also replace three vertices, one of each color.

3.3. Compounds involving the 600-cell

If Q is the 600-cell, then E is a 2-simple, 2-simplicial polytope with f -vector $(720, 3600, 3600, 720)$ and fatness exactly 5. Again, we can glue simplices onto any subset of nonadjacent facets of the 600-cell, creating convex compounds with dihedral angle $14\pi/15$. We did not count the (large) number of distinct ways of choosing such a subset, analogous to Proposition 10. It is not possible to glue a cross polytope to a 600-cell, because that would create an $11\pi/10$ angle (*i.e.*, a concave dihedral of $9\pi/10$) which cannot be filled by additional simplices or cross polytopes.

The large dihedral angles of the 600-cell make it difficult to form compounds from it, but we can modify it as follows to create smaller dihedrals. Remove a vertex and form the convex hull of the remaining 119 vertices. The resulting convex polytope has 580 of the 600-cell's tetrahedral facets and one icosahedral facet. The pentagonal edge link (Figure 4(a)) of the edges bordering this new facet become modified in a similar way, by removing one vertex and forming the convex hull of the remaining four vertices (Figure 4(b)), which results in a trapezoid; thus, the hyperbolic dihedrals at the ridges around the new facet are $2\pi/5$.

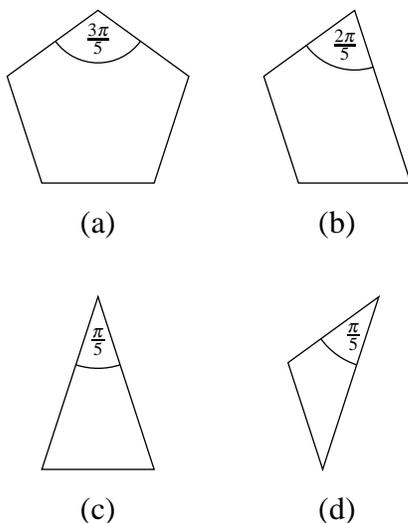


Figure 4: Edge figures of (a) a 600-cell, (b) a 600-cell with one vertex removed, (c) a 600-cell with two removed vertices, and (d) an icosahedral cap.

This cut polytope is not simplicial, but we may glue two of these polytopes together along their icosahedral facets, forming a simplicial polytope with $4\pi/5$ dihedrals along the glued ridges. This compound's new edge links are hexagons formed by gluing pairs of trapezoids (Figure 5(a)). The same cutting and gluing process may be repeated to form a sequence or tree of 600-cells, connected along cuts that do not share a ridge. For such a chain or tree formed from n cut 600-cells, the f -vector may be computed as

$$f(Q) = (106n + 14, 666n + 54, 1120n + 80, 560n + 40),$$

so the E -construction yields

$$f(E) = (666n + 54, 3360n + 240, 3360n + 240, 666n + 54)$$

and thus a fatness of

$$\phi(E) = \frac{3360n + 240}{666n + 54} \rightarrow \frac{560}{111} \approx 5.045045.$$

Thus the fatness of the 2-simple, 2-simplicial polytopes formed by such compounds improves slightly on that formed from the 600-cell alone.

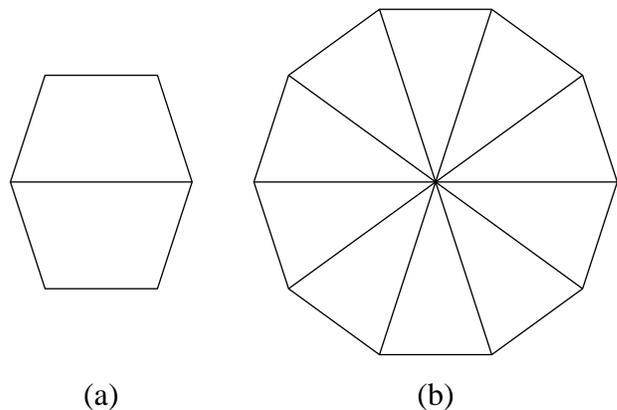


Figure 5: Edge links of compounds of cut 600-cells.

It is also possible to form compounds involving 600-cells which have been cut by removing several vertices (as described above) so that two of the resulting icosahedral facets meet at a ridge. Each edge link at this ridge is an isosceles triangle formed by removing two vertices from a pentagon (Figure 4(c)). Thus the dihedral angle of the triangular ridge between the icosahedral facets is $\pi/5$. We can therefore form compounds in which ten of these doubly-cut 600-cells share a triangle, whose edges links are a regular decagon cut into ten isosceles triangles (Figure 5(b)). Yet other compounds of cut 600-cells and simplices are possible, although we do not need them here.

The cut 600-cells also form more complicated compounds which require some group-theoretic terminology to explain. The vertices of a regular 600-cell form a 120-element group under (rescaled) quaternionic multiplication, the binary icosahedral group. This group has a 24-element subgroup, the binary tetrahedral group, which also arises as the units of the Hurwitz integers (see Conway and Sloane [7, §2.2.6,8.2.1]). Let A be the convex hull of the other 96 vertices of the 600-cell; *i.e.*, A is formed by cutting 24 vertices from 600-cell in the above manner. The resulting polytope is the “snub $\{3, 4, 3\}$ ” (snub 24-cell) of Coxeter [8, §8.4,8.5]. Its f -vector is $(96, 432, 480, 144)$. Every icosahedral facet of A is adjacent to 8 other icosahedron facets, as well as to 12 tetrahedra. Thus A has 96 icosahedron-icosahedron ridges.

We can build new hyperbolic, edge-tangent, simplicial polytopes by gluing copies of A along icosahedral faces and

capping the remaining icosahedral facets with pyramidal caps of the type that we had cut off to form A . (The edge links of such a cap C are given by Figure 4(d).) The resulting polytope Q will be convex if at each icosahedral-icosahedral ridge of a copy of A , either 10 copies of A meet, or two caps and one or two copies of A do. Also at each icosahedral-tetrahedral ridge of a copy of A , either two copies of A or one each of A and a cap must meet. If two copies of A meet (in an icosahedral facet F), then they differ by a reflection through F . These reflections generate a discrete hyperbolic reflection group Γ since the supporting hyperplanes of the icosahedral facets (the facets of a hyperideal 24-cell, whose ridges are also ridges of A !) satisfy the Coxeter condition: When they meet, they meet at an angle of $\pi/5$, which divides π . Thus the copies of A used in Q are a finite subset Σ of the orbit of A under Γ . The set Σ determines Q .

Remark. That it suffices to consider the dihedral angles of adjacent facets follows from Poincaré’s covering-space argument: Let P be a spherical, Euclidean, or hyperbolic polytope whose dihedral angles divide π . Let X denote the space in which it lives. Let Y be the disjoint union of all copies of P in X in every position, and let Z be the quotient of Y given by identifying two copies of P along a shared facet. The space Z is constructed abstractly so that P tiles it.

We claim that Z is a covering space of X . Each $p \in Z$ lies in the interior of some face F of a copy of P . If F is a copy of P or a facet, this is elementary; if F is a ridge, it follows from the dihedral angle condition. Otherwise it follows by applying the covering-space argument inductively, replacing X by the link S of F and P by $P \cap S$.

Since Z is a covering space, a connected component of Z is a tiling of X by P . See Vinberg [32] for a survey of hyperbolic reflection groups.

There is no one best choice for Σ , only a supremal limit. One reasonable choice for Σ is the corona of a copy of A , *i.e.*, A together with the set of all images under Γ that meet it, necessarily at a facet or a ridge. The corona of A is depicted by a simplified (and therefore erroneous) schematic in Figure 6; the reader should imagine the correct, more complicated version. The schematic uses a chemistry notation in which each copy of A is represented as an *atom*, each pair of copies that shares a facet is represented as a *bond*, and “ $7A$ ” denotes a chain of 7 atoms. To extend the terminology, we call 10 copies of A that meet at a ridge a *ring*. The schematic is simplified in that the central copy actually has 24 bonds (not 6), each of the neighbors has 9 bonds (not 3), and there are 96 rings (not 6).

Let Q be the union of these copies of A with the remaining icosahedral facets capped. To compute the f -vector of Q it is easier to view each atom as a copy of a 600-cell B , minus two caps for each bond. There are $1 + 24 + 96 \cdot 7 = 697$ atoms and $24 + 96 \cdot 8 = 792$ bonds in total. Thus Q has

$$f_3(Q) = 697f_3(B) - 792 \cdot 2f_3^i(C) = 386520$$

facets, where $f_3^i(C) = 30$ is the number of simplicial facets of a cap C .

Counting vertices is more complicated. Let I be an ico-

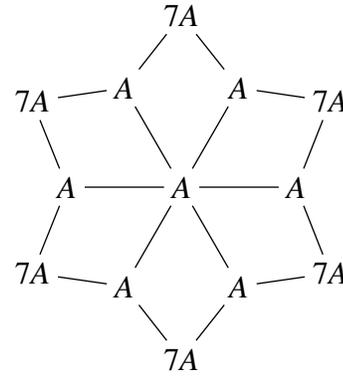


Figure 6: An oversimplified schematic of a corona of A .

hedron and let T be a triangle. Then

$$\begin{aligned} f_0(Q) &= 697f_0(B) - 792 \cdot 2f_0(C) + 792f_0(I) + 96f_0(T) \\ &= 72840, \end{aligned}$$

because after the vertices of the caps are subtracted, the vertices of each icosahedral facet at a bond are undercounted once, and after these are restored the vertices of each triangle at the center of a ring are undercounted once. The rest of the f -vector of Q follows from the Dehn-Sommerville equations:

$$\begin{aligned} f_2(Q) &= 2f_3(Q) = 773040, \\ f_1(Q) &= f_0(Q) + f_3(Q) = 459360. \end{aligned}$$

The polytope Q yields an E -polytope with fatness

$$\phi(E) = 6 \frac{f_3(Q)}{f_1(Q)} = \frac{3221}{638} \approx 5.048589.$$

Note that since this bound arises from a specific choice of Σ rather than a supremal limit, this is not optimal as a lower bound of supremal fatness.

3.4. Kissing numbers

As mentioned above, another use of ridge-tangent polytopes P is the average kissing number problem [18]. Let X be a finite packing of (not necessarily congruent) spheres S^3 , which is equivalent to a finite sphere packing in \mathbb{R}^3 by stereographic projection. The question is to maximize the average number of kissing points of the spheres in X . If P is ridge-tangent, its facets intersect the unit sphere S^3 in a sphere packing X , in which the spheres kiss at the tangency points of P . Thus the ridge-facet ratio of P is exactly half the average kissing number of X . (Not all sphere packings come from ridge-tangent polytopes in this way.)

The sphere packings due to Kuperberg and Schramm [18] can be viewed as coming from a compound consisting of a chain or tree of n cut 600-cells (*i.e.*, atoms in the sense of Figure 6). Their average kissing numbers are

$$\kappa = 2 \frac{666n + 54}{106n + 14} \rightarrow \frac{666}{53} \approx 12.56603.$$

By contrast if Q is the compound formed from a corona of A as in Section 3.3, then the average kissing number of the corresponding sphere packing is

$$\kappa(Q) = 2 \frac{f_1(Q)}{f_0(Q)} = \frac{7656}{607} \approx 12.61285.$$

Like the bound on fatness, it is not optimal as a lower bound on the supremal average kissing number.

Here we offer no improvement on the upper bound

$$\kappa < 8 + 4\sqrt{3} \approx 14.92820$$

from [18], even though it cannot be optimal either.

4. 3-SPHERES

In this section we construct a family of strongly regular cellulations of the 3-sphere with unbounded fatness. Indeed, we provide an efficient version of the construction, in the sense that it requires only polynomially many cells to achieve a given fatness. (The construction is also a polynomially effective randomized algorithm.) Given N , we find a strongly regular cellulation of S^3 with $O(N^{12})$ cells and fatness at least N . Note that there are also power-law upper bounds on fatness: An $O(N^{1/3})$ upper bound for the fatness of a convex 4-polytope with N vertices follows from work by Edelsbrunner and Sharir [9]. The Kővari-Sós-Turán theorem [17] (see also [5, p. 1239] and [23, Thm. 9.6, p. 121]) implies an $O(N^{2/3})$ upper bound on the fatness of strongly regular cellulations of S^3 , since the vertex-facet (atom-coatom) incidence graph has no $K_{3,3}$ -subgraph, and thus has at most $O(f_0(f_0 + f_3)^{2/3})$ edges. Our construction provides an $\Omega(N^{1/12})$ lower bound.

The inefficient construction is a simpler version which we describe first. For every $g > 0$, S^3 can be realized as

$$H_1 \cup (S_g \times I) \cup H_2,$$

a thickened surface of genus g capped on both ends with handlebodies. (This is obtained from a neighborhood of the standard [unknotted] smooth embedding of S_g into S^3 .) If for some g we can find a fat cellulation of S_g , we can realize S^3 as a “fat sausage with lean ends,” as shown in Figure 7. We cross the fat cellulation of S_g with an interval divided into N segments to produce a fat cellulation of $S_g \times I$. Then we fix arbitrary strongly regular cellulations of the handlebodies H_1 and H_2 . If we make the sausage long enough, *i.e.*, if we take $N \rightarrow \infty$, the fatness of the sausage converges to the fatness of its middle regardless of the structure of its ends.

It remains only to show that there are strongly regular fat cellulations of surfaces. The surface S_g has perfect cellulations with f -vector $(1, 2g, 1)$. Such a cellulation is obtained by gluing pairs of sides of a $4g$ -gon in such a way that all vertices are identified. It has fatness g , and it exists for arbitrarily large g , but it is far from regular. However, its lift to the universal cover \hat{S}_g is strongly regular, since any such cellulation can be represented by a tiling of the hyperbolic

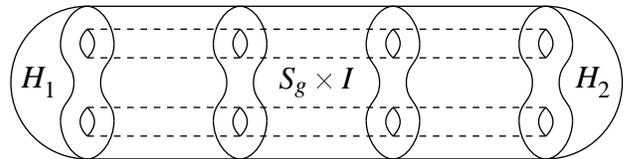


Figure 7: S^3 as a fat sausage with lean ends.

plane by *convex* polygons. (Indeed, if we take the regular $4g$ -gon with angles of $\pi/(2g)$, which is certainly convex, then its edges and angles are compatible with any perfect cellulation.) Moreover, Mal'cev's theorem [19], states that finitely generated matrix groups are residually finite; this implies that every closed hyperbolic manifold admits intermediate finite covers with arbitrarily large injectivity radius. (See [11, §4] for a detailed exposition.) In particular S_g admits an intermediate cover \hat{S}_g whose injectivity radius exceeds the diameter of a 2-cell. The cellulation of \hat{S}_g is then strongly regular. Its genus is much larger than g , but its fatness is still g .

The efficient construction is the same: It only requires careful choices for the finite cover \hat{S}_g and for the handlebodies H_1 and H_2 . Among the perfect cellulations of S_g , a convenient one for us is a $4g$ -gon with opposite edges identified. We describe the fundamental group $\pi_1(S_g)$ using this cellulation. As shown in Figure 8, we number the edges x_0, \dots, x_{4g-1} consecutively, so that

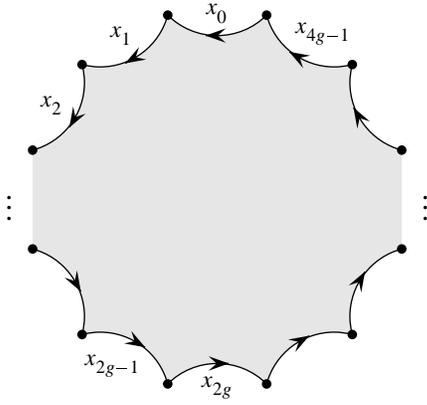
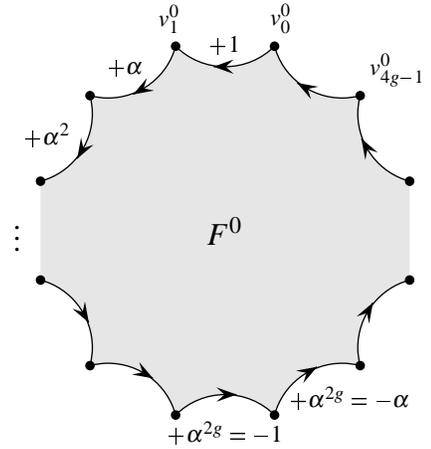
$$x_i = x_{2g+i}^{-1} \quad (1)$$

and

$$x_0 x_1 \cdots x_{4g-1} = 1. \quad (2)$$

(We interpret the indices as elements of $\mathbb{Z}/(4g)$.) Equation (1) expresses the identifications, while equation (2) expresses the boundary of the 2-cell.

We construct \hat{S}_g as a tower of two abelian finite covers, which together form an irregular cover of S_g . (No abelian cover of S_g is strongly regular.) The surface S_g satisfies the usual requirements of covering-space theory (see Fulton [12, § 13b, 14a]): It is a connected, locally path-connected, and locally simply connected space. For any finite group A , the A -coverings of S_g are covering spaces of the form Y with $Y/A = S_g$, where A acts properly discontinuously on Y . These coverings, up to isomorphism, are classified by the set of group homomorphisms $\text{Hom}(\pi_1(S_g, x), A)$ [12, Thm. 14.a]. Furthermore, if A is abelian, then every such homomorphism maps the commutators in $\pi_1(S_g, x)$ to zero, so the A -coverings are classified by $\text{Hom}(H_1(S_g, \mathbb{Z}), A)$. In other words, if A is any abelian group with n elements, then every homomorphism $\sigma : H_1(S_g) \rightarrow A$ defines an n -fold abelian covering of S_g . (If the homomorphism is not surjective, then the covering space is not connected [12, p. 193]. In this case we use a connected component of the covering space, which has the same fatness but smaller genus.)

Figure 8: Labelling edges of S_g .Figure 9: Labelling the vertices of F^0 .

Now assume that $q = 4g + 1$ is a prime power and let α generate the cyclic group \mathbb{F}_q^* , where \mathbb{F}_q is the field with q elements. We can fulfill the assumption by changing g by a bounded factor. (Most simply we can let $q = 5^k$. Or we can let q be prime, so that $\mathbb{F}_q = \mathbb{Z}/q$, by a form of Bertrand's postulate for primes in congruence classes. This result dates to the 19th century; see Erdős [10] for an elementary proof.) Define a homomorphism $\sigma : H_1(S_g) \rightarrow \mathbb{F}_q$ by $\sigma([x_i]) = \alpha^i$, where $[x_i]$ is the 1-cycle (or homology class) represented by the loop x_i . Since $\alpha^{2g} = -1$, the definition of σ is consistent with equation (1). Consistency with (2) is then automatic. Let S_g^l be the finite cover corresponding to σ .

To prepare for the subsequent analysis of \widehat{S}_g , we give an explicit combinatorial description of S_g^l . Let F^0 be a 2-cell of S_g^l and label its vertices

$$v_0^0, v_1^0, \dots, v_{4g-1}^0$$

in cyclic order. See Figure 9.

For each $s \in \mathbb{F}_q$, let F^s and v_k^s be the images of F^0 and v_k^0 under the action of s . Since S_g has only one vertex, the vertices of S_g^l may be identified with \mathbb{F}_q . Thus, if we identify v_0^0 with $0 \in \mathbb{F}_q$, then the action of \mathbb{F}_q will identify v_0^s with $s \in \mathbb{F}_q$. The structure of σ further implies that

$$v_k^s = s + 1 + \alpha + \alpha^2 + \dots + \alpha^{k-1} = s + \frac{\alpha^k - 1}{\alpha - 1}$$

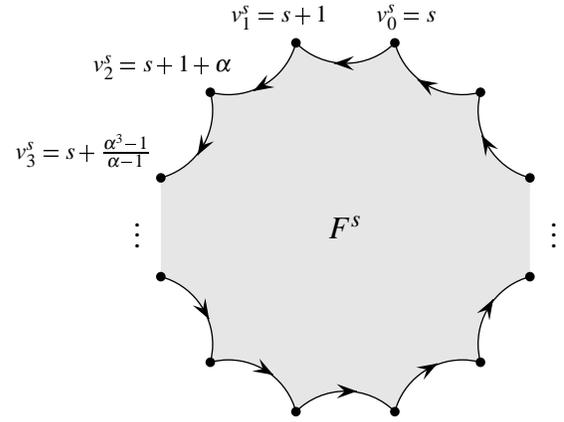
for all $k \in \mathbb{Z}/(4g)$ and $s \in \mathbb{F}_q$. See Figure 10.

Using this explicit description, it is routine to verify the following (remarkable) properties of the surface S_g^l .

Lemma 12. *The cellulation of the abelian cover S_g^l is regular and has f -vector*

$$f(S_g^l) = (q, 2gq, q) = (q, \binom{q}{2}, q).$$

Every facet has $q - 1 = 4g$ vertices, while every vertex has degree $q - 1 = 4g$. The graph (or 1-skeleton) of S_g^l is the complete graph on $q + 1$ vertices. The dual graph is also complete;

Figure 10: Labelling the vertices of F^s .

any two facets share exactly one edge (as well as $q - 4$ other vertices).

Proof. In view of the combinatorial description above (Figure 10), all these facts follow from simple computations in the field \mathbb{F}_q :

- S_g^l is regular — for each $s \in \mathbb{F}_q$, the vertex labels $s + \frac{\alpha^k - 1}{\alpha - 1}$ ($0 \leq k < 4g$) are distinct.
- The 1-skeleton of S_g^l is complete — for $v, v' \in \mathbb{F}_q$, $v \neq v'$ there is a unique $s \in \mathbb{F}_q$ and $k \in \mathbb{Z}/(4g)$ with

$$v = s + \frac{\alpha^k - 1}{\alpha - 1} \quad v' = s + \frac{\alpha^{k+1} - 1}{\alpha - 1}.$$

- The dual graph of S_g^l is complete — for $s, s' \in \mathbb{F}_q$, $s \neq s'$,

there are unique $k, \ell \in \mathbb{Z}/(4g)$ such that

$$\begin{aligned} s + \frac{\alpha^k - 1}{\alpha - 1} &= s' + \frac{\alpha^{\ell+1} - 1}{\alpha - 1} \\ s + \frac{\alpha^{k+1} - 1}{\alpha - 1} &= s' + \frac{\alpha^\ell - 1}{\alpha - 1}. \end{aligned}$$

□

Theorem 13. *Let $n \geq 128g^4$, and let*

$$\rho : H_1(S'_g) \rightarrow \mathbb{Z}/n$$

be a randomly chosen homomorphism, and let \widehat{S}_g be the finite cover of S'_g corresponding to ρ . Then with probability more than $\frac{1}{2}$, the cellulation of \widehat{S}_g is strongly regular.

In order to prove Theorem 13, we need to more explicitly describe the condition of strong regularity as it applies to \widehat{S}_g . Let X be a regular cell complex and suppose that its universal cover \widetilde{X} is strongly regular. Recall that the *star* $\text{st}(v)$ of a vertex v in X is the subcomplex generated by the cells that contain v . The complex X is strongly regular if and only if the star of each vertex is. Suppose that $\tilde{v} \in \widetilde{X}$ projects to $v \in X$. Then the star $\text{st}(\tilde{v})$, which is strongly regular, projects to the star $\text{st}(v)$. The latter is strongly regular if and only if the projection is injective. In other words, X is strongly regular if and only if the stars of \widetilde{X} embed in X . If X is not strongly regular, then there must be a path $\tilde{\ell}$ connecting distinct vertices of $\text{st}(\tilde{v})$ which projects to a loop ℓ in $\text{st}(v)$. We say that such a loop *obstructs strong regularity*. We can assume that $\tilde{\ell}$ is a pair of segments properly embedded in distinct cells in $\text{st}(\tilde{v})$, with only the end-points of the segments on the boundary of the cells, which implies that ℓ is embedded if X is regular. Figure 11 gives an example of such a loop ℓ in a regular cellulation of a torus.

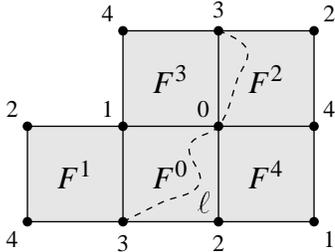


Figure 11: A loop ℓ that obstructs strong regularity in the torus S'_1 .

In our case, the surfaces S'_g are regular, but they have many obstructing loops. Theorem 13 asserts that, with non-zero probability, all such loops lengthen when lifted to \widehat{S}_g .

Lemma 14. *No loops in S'_g that obstruct strong regularity are null-homologous. Furthermore, all obstructing loops represent indivisible elements in $H_1(S'_g)$.*

Proof. In brief, they are indivisible because they are embedded, and they are too short to be null-homologous.

By Lemma 12, S'_g is regular. By the discussion after the statement of Theorem 13, each obstructing loop ℓ is embedded. If ℓ separates S'_g , then it is null-homologous. If ℓ does not separate S'_g , then it is indivisible in homology. (To show this, we can appeal to the classification of surfaces by cutting S'_g along ℓ . The classification implies that all non-separating positions for ℓ are equivalent up to homeomorphism of S'_g . It is easy to find a standard position for ℓ in which it is indivisible in homology.) Thus it remains to show that no obstructing loop is null-homologous.

First, we claim that any obstructing loop ℓ can be supported on fewer than $4g$ edges of the 1-skeleton of S'_g . We homotop the two segments of ℓ to the boundaries of the 2-cell containing them, giving them each at most $2g$ edges. Thus ℓ is represented by a sequence of at most $4g$ edges in S'_g . The case of exactly $4g$ edges does not occur, since the endpoints of the loop coincide, and no two vertices $v \neq v'$ of S'_g are opposite vertices in two different facets F^s . As in Lemma 12, this follows from the fact that for $v, v' \in \mathbb{F}_g$, $v \neq v'$, there are unique $s \in \mathbb{F}_q$ and $k \in \mathbb{Z}/(4g)$ with

$$v = s + \frac{\alpha^k - 1}{\alpha - 1} \quad v' = s + \frac{\alpha^{k+2g} - 1}{\alpha - 1}.$$

Second, we claim that any null-homologous loop in the 1-skeleton of S'_g contains at least $4g$ edges. In other words, if f is a 2-chain on S'_g and $\partial f \neq 0$, then $|\partial f| \geq 4g$. Since S'_g is orientable, we can regard f as a function on its 2-cells. Since f is non-constant, it attains some value t on k 2-cells with $0 < k < q$. Since any two 2-cells share an edge by Lemma 12, these 2-cells share

$$k(q-k) \geq 4g$$

edges with the complementary set of 2-cells, of which there are $q-k$. Since ∂f is non-zero on these edges, $|\partial f| \geq 4g$, as desired. □

Proof of Theorem 13. In brief, S'_g has fewer than $64g^4$ obstructing loops ℓ . For each one,

$$P[\rho([\ell]) = 0] = \frac{1}{n}.$$

The expected number of obstructing loops that lift from S'_g to \widehat{S}_g without lengthening is less than $64g^4/n \leq \frac{1}{2}$. Thus there is a good chance that all obstructing loops lengthen.

The homology group $H_1(S'_g)$ is a finitely generated free abelian group: It is isomorphic to \mathbb{Z}^d , for $d = 1 + q(g-1)$. Thus it admits n^d homomorphisms ρ to \mathbb{Z}/n . Since this is a finite number, choosing one uniformly at random is well-defined. If c is any indivisible vector in \mathbb{Z}^d , then it is contained in a basis, and thus $\rho(c)$ is equidistributed. In particular, if ℓ is an obstructing loop, then $[\ell]$ is indivisible by Lemma 14, so $\rho([\ell])$ is equidistributed in \mathbb{Z}/n .

It remains only to bound the number of obstructing loops in S'_g . A star $\text{st}(v)$ in \widehat{S}_g has $4g(4g-2)$ points other than v itself.

Without loss of generality, v projects to 0 in S'_g . In this case the other vertices are equidistributed among the $4g$ non-zero elements of \mathbb{F}_q . Therefore $\text{st}(v)$ has

$$4g \binom{4g-2}{2} = \frac{4g(4g-2)(4g-3)}{2}$$

pairs of arcs connecting v to two vertices that are the same in S'_g . These pairs represent all two-segment obstructing loops that pass through 0 and a nonzero vertex v' (and some of these loops are homotopic). If we count such pairs of arcs for any pair of distinct vertices v, v' of S'_g , then we find that the total number is not more than

$$\binom{4g+1}{2} \binom{4g-2}{2} = \frac{(4g+1)4g(4g-2)(4g-3)}{4} < 64g^4,$$

as desired. \square

Question 15. For each $g > 1$, what is the maximum fatness of a strongly regular cellulation of a surface of genus g ? Equivalently, how many edges are needed for a strongly regular cellulation of a surface of genus g ?

Remark. One interesting alternative to the construction of S'_g is to assume instead that $q = 4g - 1$ is a prime power, and to let α be an element of order $4g$ in \mathbb{F}_{q^2} . The resulting q^2 -fold cover S''_g is almost strongly regular: the only obstructing loops are those that are null homologous in S_g . Another interesting surface is the modular curve $X(2p)$, where p is a prime [25, §13]. The inclusion $\Gamma(2p) \subset \Gamma(2)$ of modular groups induces a projection from $X(2p)$ to the modular curve $X(2)$, which is a sphere with three cusp points. If we connect two of these points by an arc which avoids the third, it lifts to a cellulation of $X(2p)$ with f -vector $(p^2 - 1, \frac{p(p^2-1)}{2}, \frac{p^2-1}{2})$. Like S''_g , it has a few obstructing loops. Unfortunately we do not know a way to use either S''_g or $X(2p)$ to make fat surfaces of lower genus (or equivalently fewer cells) than \widehat{S}_g .

Since Theorem 13 provides us with efficient fat surfaces \widehat{S}_g , the construction of fat cellulations of S^3 only requires efficient cellulations of the handlebodies H_1 and H_2 and an efficient way to attach them to \widehat{S}_g . In our construction the handlebody cellulations are *a priori* unrelated to the cellulation of \widehat{S}_g . Rather they are transverse after attachment, and each point of intersection will become a new vertex. Thus the question is to position the cellulations to minimize their intersection.

We describe the cellulations in three stages: first, a dissection of H_1 and H_2 individually into 3-cells; second, their relative position; and third, their position relative to the cellulation of \widehat{S}_g .

Let \widehat{g} be the genus of \widehat{S}_g . A handlebody H of genus \widehat{g} can be formed by identifying \widehat{g} pairs of disks on the surface of a 3-cell. The result is a dissection of H into \widehat{g} 2-cells and one 3-cell, although it is not a cell complex because there are no

1-cells or 0-cells. We can still ask whether such a dissection is regular or strongly regular; this one is neither. However, if we replace each 2-cell by 3 parallel 2-cells, it becomes strongly

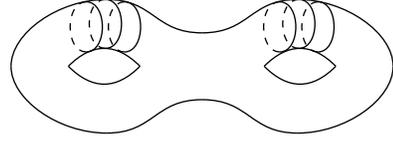


Figure 12: A strongly regular dissection A of a handlebody of genus 2.

regular. An example of the resulting dissection A is shown in Figure 12.

The surface $S_{\widehat{g}}$ (which for our choice of \widehat{g} is isomorphic to \widehat{S}_g) has another standard perfect cellulation called a *canonical schema* in the computer science literature [31]. Using the labelling in Figure 8, we identify x_{4k} with x_{4k+2}^{-1} (the even loops), and x_{4k+1} with x_{4k+3}^{-1} (the odd loops), for all $0 \leq k \leq \widehat{g}$. By rounding corners we can make each even loop intersect one odd loop once and eliminate all other intersections between loops. If we interpret this pattern of loops as the standard Heegaard diagram for S^3 [29], then the even loops bound disks in H_1 and the odd loops bound disks in H_2 . We can then put in two copies A_1 and A_2 of the cell division A so that its 2-cells run parallel to these loops.

We would like to position A_1 and A_2 to minimize their intersection with the cellulation of \widehat{S}_g . (Note that A_1 and A_2 do not intersect each other since $\widehat{S}_g \times I$ lies in between.) To this end Vegter and Yap [31] proved that any cellulation of $S_{\widehat{g}}$ with n edges admits a position of the canonical schema cellulation with $O(n\widehat{g})$ intersections. (Strictly speaking the theorem applies to triangulations, but any regular cellulation of a surface with n edges can be refined to a triangulation with less than $3n$ edges.) In our case

$$n = \Theta(\widehat{g}) = \Theta(g^6).$$

Thus, by tripling the edges of the canonical schema in the Vegter-Yap construction, we can position A_1 and A_2 so that the lean ends in the sausage have $O(g^{12})$ vertices. If we give the fat part of the sausage $N = g^{12}$ slices, the total f -vector of the cellulation of S^3 is then

$$(\Theta(g^{12}), \Theta(g^{13}), \Theta(g^{13}), \Theta(g^{12})),$$

and its fatness is $\Theta(g)$. This completes the efficient construction with unbounded fatness.

-
- [1] David Barnette, *Inequalities for f -vectors of 4-polytopes*, Israel J. Math. **11** (1972), 284–291.
- [2] Margaret M. Bayer, *The extended f -vectors of 4-polytopes*, J. Combinatorial Theory Ser. A **44** (1987), no. 1, 141–151.
- [3] Margaret M. Bayer and Louis J. Billera, *Generalized Dehn-Sommerville relations for polytopes, spheres and Eulerian partially ordered sets*, Inventiones Math. **79** (1985), no. 1, 143–157.
- [4] Louis J. Billera and Carl W. Lee, *A proof of the sufficiency of McMullen's conditions for f -vectors of simplicial convex polytopes*, J. Combin. Theory Ser. A **31** (1981), no. 3, 237–255.
- [5] Béla Bollobás, *Extremal graph theory*, Handbook of Combinatorics (R. Graham, M. Grötschel, and L. Lovász, eds.), North-Holland/Elsevier, Amsterdam, 1995, pp. 1231–1292.
- [6] Tom Braden, *A glued hypersimplex*, 1997, Personal communication.
- [7] John H. Conway and Neil J. A. Sloane, *Sphere packings, lattices and groups*, 3rd ed., Grundlehren der mathematischen Wissenschaften, vol. 290, Springer-Verlag, New York, 1993.
- [8] Harold Scott MacDonald Coxeter, *Regular polytopes*, 2nd ed., Macmillan, New York, 1963, Corrected reprint, Dover, New York 1973.
- [9] Herbert Edelsbrunner and Micha Sharir, *A hyperplane incidence problem with applications to counting distances*, Applied geometry and discrete mathematics: The Victor Klee Festschrift (P. Gritzman and B. Sturmfels, eds.), DIMACS Series in Discrete Mathematics and Theoretical Computer Science, vol. 4, Amer. Math. Soc., Providence, RI, 1991, pp. 253–263.
- [10] Paul Erdős, *Über die Primzahlen gewisser arithmetischer Reihen*, Math. Z. **39** (1935), 473–491.
- [11] Gábor Fejes Tóth, Greg Kuperberg, and Włodzimierz Kuperberg, *Highly saturated packings and reduced coverings*, Monatsh. Math. **125** (1998), no. 2, 127–145.
- [12] William Fulton, *Algebraic topology. A first course*, Graduate Texts in Mathematics, vol. 153, Springer-Verlag, New York, 1995.
- [13] Branko Grünbaum, *Convex polytopes*, Interscience, London, 1967.
- [14] Andrea Höppner and Günter M. Ziegler, *A census of flag-vectors of 4-polytopes*, Polytopes — combinatorics and computation (G. Kalai and G.M. Ziegler, eds.), DMV Seminars, vol. 29, Birkhäuser-Verlag, Basel, 2000, pp. 105–110.
- [15] Birger Iversen, *Hyperbolic geometry*, London Math. Soc. Student Texts, vol. 25, Cambridge University Press, Cambridge, 1992.
- [16] Michael Joswig and Günter M. Ziegler, *Neighborly cubical polytopes*, Discrete Comput. Geom. **24** (2000), no. 2-3, 325–344, arXiv:math.CO/9812033.
- [17] Tamás Kővári, Vera T. Sós, and Pál Turán, *On a problem of K. Zarankiewicz*, Colloq. Math. **3** (1954), 50–57.
- [18] Greg Kuperberg and Oded Schramm, *Average kissing numbers for non-congruent sphere packings*, Math. Res. Lett. **1** (1994), no. 3, 339–344, arXiv:math.MG/9405218.
- [19] A. I. Mal'cev, *On the faithful representation of infinite groups by matrices*, Amer. Math. Soc. Transl. (2) **45** (1965), 1–18.
- [20] Joseph Malkevitch, *Tiling convex polygons with equilateral triangles and squares*, Discrete Geometry and Convexity (J. E. Goodman, E. Lutwak, J. Malkevitch, and R. Pollack, eds.), Ann. New York Acad. Sci., vol. 440, New York Academy of Sciences, 1985, pp. 299–303.
- [21] Peter McMullen, *The numbers of faces of simplicial polytopes*, Israel J. Math. **9** (1971), 559–570.
- [22] James R. Munkres, *Elements of algebraic topology*, Addison-Wesley, Menlo Park, CA, 1984.
- [23] János Pach and Pankaj K. Agarwal, *Combinatorial geometry*, J. Wiley and Sons, New York, 1995.
- [24] Egon Schulte, *Analogues of Steinitz's theorem about non-inscribable polytopes*, Intuitive geometry (Siófok 1985) (Amsterdam), Colloquia Soc. János Bolyai, vol. 48, North Holland, 1987, pp. 503–516.
- [25] Joseph H. Silverman, *The arithmetic of elliptic curves*, Graduate Texts in Mathematics, vol. 106, Springer-Verlag, New York, 1986.
- [26] Richard P. Stanley, *The number of faces of a simplicial convex polytope*, Adv. in Math. **35** (1980), no. 3, 236–238.
- [27] ———, *A survey of Eulerian posets*, Polytopes: abstract, convex and computational (Scarborough, ON, 1993), NATO Adv. Sci. Inst. Ser. C Math. Phys. Sci., vol. 440, Kluwer Acad. Publ., Dordrecht, 1994, pp. 301–333.
- [28] Ernst Steinitz, *Über die Eulerschen Polyederrelationen*, Archiv für Mathematik und Physik **11** (1906), 86–88.
- [29] John Stillwell, *Classical topology and combinatorial group theory*, 2nd ed., Graduate Texts in Mathematics, no. 72, Springer-Verlag, 1993.
- [30] William P. Thurston, *Three-dimensional geometry and topology*, vol. I, Princeton University Press, Princeton, NJ, 1997.
- [31] Gert Vegter and Chee K. Yap, *Computational complexity of combinatorial surfaces*, Proc. 6th ACM Symp. Comp. Geom., ACM Press, 1990, pp. 102–111.
- [32] Ernest B. Vinberg, *Hyperbolic reflection groups*, Russian Math. Surveys **40** (1985), 31–75.
- [33] William Waite, *Combining squares and triangles*, Cubism for Fun **53** (2000), 18–22.
- [34] David W. Walkup, *The lower bound conjecture for 3- and 4-manifolds*, Acta Math. **125** (1970), 75–107.
- [35] Günter M. Ziegler, *Lectures on polytopes*, Graduate Texts in Mathematics, vol. 152, Springer-Verlag, New York, 1995, Revised edition, 1998.