Extending a triangulation from the 2-sphere to the 3-ball

Peter Doyle Matthew Ellison Zili Wang

Version 1.000 dated 19 March 2023 No Copyright^{*}

Abstract

Define the tet-volume of a triangulation of the 2-sphere to be the minimum number of tetrahedra needed to extend it to a triangulation of the 3-ball, and let d(v) be the maximum tet-volume for v-vertex triangulations. In 1986 Sleator, Tarjan, and Thurston (STT) proved that d(v) = 2v - 10 holds for large v, and conjectured that it holds for all $v \ge 13$. Their proof used hyperbolic polyhedra of large volume. They suggested using more general notions of volume instead, and Mathieu and Thurston showed the potential of this approach in a paper that has been all but lost. Taking this as our cue, we prove the conjecture. This implies STT's associated conjecture, proven by Pournin in 2014, about the maximum rotation distance between trees.

For Bill

1 Summary

A triation of the sphere is an oriented simplicial 2-complex σ whose carrier is the 2-sphere. Regard σ as a subcomplex of the (v-1)-simplex Δ^{v-1} , where v is the number of vertices. A tetration of σ is an oriented 3-subcomplex τ of Δ^{v-1} with boundary $\partial \tau = \sigma$. Define the tet-volume tetvol(σ) to be

^{*}The authors hereby waive all copyright and related or neighboring rights to this work, and dedicate it to the public domain. This applies worldwide.

the minimum number of 'tets' (tetrahedra) in a tetration. Let d(v) be the maximum of $\mathbf{tetvol}(\sigma)$ over all v-vertex triations of the 2-sphere.

In 1986 Sleator, Tarjan, and Thurston (STT) observed [3, 4] that if $v \ge 13$ then $d(v) \le 2v - 10$, as follows. From Euler, a triation with v vertices has 2v - 4 faces. Coning from a vertex to non-adjacent faces yields a tetration, so

 $\mathbf{tetvol}(\sigma) \le 2v - 4 - \mathbf{maxdeg}(\sigma).$

Assuming $v \ge 13$ we have $\mathbf{maxdeg}(\sigma) \ge 6$, yielding

$$\mathbf{tetvol}(\sigma) \le 2v - 10.$$

STT conjectured:

Conjecture (tetvol).

$$d(v) = 2v - 10, v \ge 13.$$

They proved this with 13 replaced by some unspecified constant, by considering triations arising as the boundary of certain ideal hyperbolic polyhedra with large volume. Because there is an upper bound for the volume of a hyperbolic tet, to tetrate a hyperbolic polyhedron with large volume requires a large number of tets.

STT suggested proving lower bounds by using more general notions of volume. Mathieu and Thurston (MT) pursued this in [1]. (See section 4.) We take this same approach here, and prove the tetvol conjecture by examining a very simple class of triations, obtained by truncating skinny cylindrical quotients of the Eisenstein lattice.

The work of STT on tetrations was motivated by the problem of finding the maximum rotation distance between trees, or equivalently, the maximum flip distance d'(v) between two triangulations of a v-gon. (See [4] for definitions and discussion.) The tetvol conjecture implies the associated conjecture that d'(v) = 2v - 10 for $v \ge 13$. (See section 7.) In 2014 Pournin [2] proved this associated conjecture directly, without reference to tet-volume, by producing pairs that he could show have flip distance 2v - 10. Pournin's theorem doesn't settle the tetvol conjecture, because there may be a gap between tet-volume and flip distance—but it removes the original motivation for studying it. Hopefully you consider the tetvol conjecture interesting in its own right. Even if you don't, the examples described here provide myriad pairs maximizing flip distance, with Pournin's pairs among them. Here's an outline of the paper. First we'll introduce the key idea of a volume potential, as applied to the icosahedron, a toy case that shows the basic idea. Then we'll apply the method to our skinny triations to settle the tetvol conjecture. We'll then discuss our debt to MT; the connection to linear programming; related constructions of triations; implications for flip distance; and where we go from here.

2 Warming up with the icosahedron

We'll begin with the icosahedron (v = 12; f = 20; maxdeg = 5). Coning from a vertex yields a tetration with 20 - 5 = 15 = 2v - 9 tets, so $tetvol(icos) \le 15$. Now let's show that $tetvol(icos) \ge 15$.

We will call a function $\rho(ABC)$ on ordered triples of vertices ABC satisfying

$$\rho(ABC) = \rho(BCA) = -\rho(ACB)$$

a volume potential. For any permutation XYZ we get from this that $\rho(XYZ) = \pm \rho(ABC)$, depending on the sign of the permutation. (In standard language, ρ is a 2-cochain.)

Associated to a volume potential is its volume form $\operatorname{vol}_{\rho}(ABCD)$, the function on ordered triples ABCD given by

$$\mathbf{vol}_{\rho}(ABCD) = \rho(BCD) - \rho(ACD) + \rho(ABD) - \rho(ABC).$$

For any permutation XYZW of ABCD we have $\mathbf{vol}_{\rho}(XYZW) = \pm \mathbf{vol}_{\rho}(ABCD)$. (In standard language, \mathbf{vol}_{ρ} is 3-cocycle, the coboundary of ρ .)

Let $\rho(\mathbf{icos})$ be the sum of $\rho(ABC)$ over the faces ABC of \mathbf{icos} . (Here and hereafter, by 'faces' we mean properly oriented faces.) For τ a tetration of \mathbf{icos} let $\mathbf{vol}_{\rho}(\tau)$ be the sum of $\mathbf{vol}_{\rho}(ABCD)$ over the tets of τ . The key fact we need is that

$$\operatorname{vol}_{\rho}(\tau) = \rho(\operatorname{icos}).$$

This is Stokes's theorem in combinatorial form; it's true because matching faces of the tets of τ make contributions of opposite sign to $\mathbf{vol}_{\rho}(\tau)$, so that after cancellation only the contributions from faces of **icos** remain.

Call the volume potential ρ good if the volume form \mathbf{vol}_{ρ} assigns all tets ABCD volume at most 1:

$$\operatorname{vol}_{\rho}(ABCD) \le 1.$$

Observe that this implies that

$$|\mathbf{vol}_{\rho}(ABCD)| \leq 1,$$

since

$$\operatorname{vol}_{\rho}(ABCD) = -\operatorname{vol}_{\rho}(ABDC) \ge -1.$$

For a good volume potential ρ the number $|\tau|$ of tets of τ satisfies

$$|\tau| \ge \operatorname{vol}_{\rho}(\tau) = \rho(\operatorname{icos}).$$

So to prove that tetvol(icos) = 15, we just need to find a good volume potential ρ with $\rho(icos) = 15$. (Or at least with $\rho(icos) > 14$, because we can always round a non-integral lower bound up: see section 5.)

Let's look for a volume potential ρ that is invariant under orientationpreserving symmetries of **icos**. (This won't hold us back: If there's any ρ at all, we can average to get a symmetrical one.)

To define ρ , we need to prescribe a value for every triple ABC of distinct vertices. We distinguish three cases.

- ABC or its orientation-reversal ACB is a face of icos. Since we want ρ to be symmetric and $\rho(\mathbf{icos}) = 15$, we must take $\rho(ABC) = 3/4$ if ABC is a face, which makes $\rho(ABC) = -3/4$ if ACB is a face.
- $\{B, C\}$ is an edge of **icos**, but not both $\{A, B\}$ and $\{A, C\}$. This is the crucial case. By symmetry, we can assume that A is any fixed vertex of **icos**. Figure 1 shows the values for a fixed choice of A by means of a flow ϕ along the edges of the dual graph. Along the dual edge BC^{\perp} clockwise from BC the flow rate is $\phi(BC^{\perp}) = \rho(ABC)$.
- ABC involves no edge of **icos**. Here we'll take $\rho(ABC) = 0$. Up to symmetry of **icos** there is just one possibility for the unoriented triangle with vertices $\{A, B, C\}$ —or two possibilities if you don't allow orientation-reversing symmetries—but for this proof we don't need to check this.

Now we want to check that ρ is good, i.e., that $\operatorname{vol}_{\rho}(ABCD) \leq 1$ for any 4-tuple ABCD. Suppose ABCD contains a face of **icos**. We may assume that BCD is a properly oriented face, with A having been moved to the standard position. We have cooked up the flow ϕ so that the net flow into



Figure 1: Volume potential flow for **icos**. Flow rates are one quarter of what is shown. To an edge BC we assign a flow of rate $\rho(ABC)$ along the dual edge BC^{\perp} . This flow is not conservative: There is net flow 3/4 out of each of the five faces incident with A, and net flow 1/4 into each of the fifteen remaining faces.

any face of **icos** not incident with A is 1/4. In terms of ρ the net flow into BCD is

$$1/4 = -\phi(BC^{\perp}) - \phi(CD^{\perp}) - \phi(DB^{\perp})$$

= $-\rho(ABC) - \rho(ACD) - \rho(ADB)$
= $-\rho(ABC) - \rho(ACD) + \rho(ABD).$

Since $\rho(BCD) = 3/4$, this gives us

$$\operatorname{vol}_{\rho}(ABCD) = \rho(BCD) - \rho(ACD) + \rho(ABD) - \rho(ABC) = 3/4 + 1/4 = 1.$$

We're now close to having shown that ρ is good. We just need to check the case where the tet ABCD contains no face of **icos**, whether properly or improperly oriented. But for such tets the four terms of $\mathbf{vol}_{\rho}(ABCD)$ all have absolute value at most 1/4, so $|\mathbf{vol}_{\rho}(ABCD)| \leq 1$.

Having verified that ρ is good, with $\rho(\mathbf{icos}) = 15$, we're done.

```
#!/usr/bin/python3
from rho import rho
# define phyllohedra T_v
def Tbase(v):
    return [(0,5,4),(0,4,3),(0,3,2),(0,2,1)]
def Tsides(v):
    pointup=list((k,k+1,k+6) for k in range(v-6))
    pointup=list((k,k+6,k+5) for k in range(v-6))
    return pointup+pointdown
def Tlid(v):
    return [(v-1,v-6,v-5),(v-1,v-5,v-4),(v-1,v-4,v-3),(v-1,v-3,v-2)]
def T(v): return Tbase(v)+Tsides(v)+Tlid(v)
# we need only check rho(T_v)=2v-10 for two values of v
def rhovol(poly): return sum(rho(abc) for abc in poly)
assert rhovol(T(13))==2*13-10 and rhovol(T(14))==2*14-10
```

Figure 2: Code to define T_v and check that $\rho(T_v) = 2v - 10$.

3 Proof of the tetvol conjecture

To prove the tetvol conjecture, all we need is a sequence of triations T_v and good volume potentials ρ_v with $\rho_v(T_v) = 2v - 10$ for all $v \ge 13$. These are defined and checked by the code in Figures 2 and 7, which together constitute a proof of the tetvol conjecture. Or rather, a 'verification'. Where's the proof?

To understand the family T_v , let's start by looking at some pictures: Figures 3 and 4. (Better yet, build some physical models!) The vertices of T_v are labeled $0, \ldots, v - 1$ spiraling up from the bottom. Cutting along edges 01 and (v-2)(v-1) we get a topological cylinder which unwraps to give a diagram like those shown in Figure 5 for v = 13 and v = 14.

These triations T_v are examples of what we call 'phyllohedra'. They are obtained as follows. Associated to the Eisenstein integers $\mathbf{Eis} = \{u + v\omega\},\$



Figure 3: Triation T_{23}



Figure 4: Triations T_{13} through T_{22}

 $\omega = \exp(i\tau/3)$ is a triation of the plane with six triangles meeting per vertex. This triation descends to the quotient cylinder

$$\Phi_{a,b} = \mathbf{Eis}/((a - b\omega)\mathbf{Z}).$$

Wrapping $\Phi_{a,b}$ around the z-axis, we get pictures like those that arise in phyllotaxis. (See Figure 6.) This inspires us to call $\Phi_{a,b}$ the (a,b)-phyllocylinder.

We can truncate a phyllocylinder by taking a subset of the vertices and the triangles they inherit from $\Phi_{a,b}$, together with some extra edges and faces to cap off the bottom and top. We call these finite triations 'phyllohedra', a loose term whose precise meaning will depend on what kinds of truncation and capping you allow.

Our triations T_v are (5, 1)-phylohedra, obtained by truncating $\Phi_{5,1}$ and then capping in the most natural way.

To accompany our T_v , we need volume potentials ρ_v . We will take these to be restrictions of a single translation-invariant volume potential ρ defined on the infinite cylinder $\Phi_{5,1}$. This is possible because the vertices of T_v are a subset of the vertices of $\Phi_{5,1}$. On any $\Phi_{a,1}$ the vertices are nicely indexed by



Figure 5: T_{13} and T_{14} unwrapped



Figure 6: The phyllocylinders $\Phi_{6,0}, \Phi_{5,1}, \Phi_{4,2}, \Phi_{4,2}, \Phi_{5,2}$. Unless *a* or *b* vanishes, $\Phi_{a,b}$ has *b* spirals in direction 1 (black); *a* spirals in direction ω (red); a + b spirals in direction $1 + \omega$ (blue).

```
#!/usr/bin/python3
def rho(abc):
  (a,b,c)=abc
  # make sure a<=b<=c</pre>
  if a>b: return -rho((b,a,c))
  if b>c: return -rho((a,c,b))
  # rho depends only on the gaps
  (x,y)=(b-a,c-b)
  if x>5:
   return 0
  elif 3 \le x \le 5 and x + y \ge 6:
    return -1
  elif x=2 and y==1:
    return 1
  elif x==1 and b>=1:
    return 1
  else:
    return 0
def vol(abcd):
   (a,b,c,d)=abcd
   return rho((b,c,d))-rho((a,c,d))+rho((a,b,d))-rho((a,b,c))
# check goodness of rho for 4-tuples of integers between 0 and n-1
# n=18 should do it, but let's go way overboard
n=36
tuples=list((a,b,c,d) for a in range(n) for b in range(n)
                      for c in range(n) for d in range(n))
vols=list(vol(abcd) for abcd in tuples)
# for this rho the only values we should see are 0,1,-1
assert set(vols) == {0,1,-1}
```

Figure 7: Code to define ρ and check that it is good.

integers, so we can think of ρ as defined for triples of integers. Translation invariance means that $\rho((a, b, c)) = \rho((a + k, b + k, c + k))$.

Figure 7 shows code to compute ρ , and check the volume condition for all 4-tuples of integers between 0 and n-1. Because of the way $\rho((a, b, c))$ depends only on the gaps between a, b, c, and treats gaps that are 6 or bigger as equal, taking n = 18 should cover all possible cases; we take n = 36 in case n = 18 is off by one (or two, or three, or four,...).

To check the goodness of ρ by hand, we can look at the volume potential flow, as we did for the icosahedron. By the translation invariance of ρ , here again we need only a single picture: Figure 8. This time the inflow vanishes for triangles other than those containing the reference vertex A, here represented by a black dot. This makes $\mathbf{vol}_{\rho}(ABCD) = 1$ whenever BCD is a face of $\Phi_{5,1}$. This takes us a long way toward showing that ρ is



Figure 8: Volume potential flow for $\Phi_{5,1}$. Reference vertex A is the black dot. Flow along red arrows has rate 1. The net flow into triangles not incident with A is 0.

good. We still have to deal with tets ABCD not involving any face of **icos**, and hence with values ρ of triangles not containing any edge. We don't know any really clever way to check these 'big' tets. Fortunately lots of triangles have $\rho = 0$, which makes things easier. We can either roll up our sleeves and get to work, or decide to trust the computer on this.

To complete the proof, we must check that $\rho(T_v) = 2v - 10$. Of the 2v - 4'side' faces of T_v , all but eight are faces of $\Phi_{5,1}$, and thus get weight 1, which gets us up to 2v - 12. We just need to check that the net contribution of the base and lid together give us the extra 2 we need. Since these faces all contain an edge of $\Phi_{5,1}$ (either 01 or (v - 2)(v - 1)) the information needed to check this is there in the flow diagram 8. As an alternative, or as a check on our work, we can observe that just have to check a single value of v to nail down the constant term, and the code in Figure 2 has done this for us.

This completes the proof.

We should emphasize that this ρ is not canonical, not even at the level of the associated volume form \mathbf{vol}_{ρ} . It's particularly nice in that it takes only values 0, 1, -1. In fact you can get the exact value of **tetvol** for small triations with such 'binary' volume potentials. But once you get up to v = 19or so, you find triations with non-integral **tetvol**, and then the jig is up.

4 Volume potentials

The volume potential method we're using here was proposed (in an equivalent form) by STT; Mathieu and Thurston (MT) explored it in [1], and introduced the flows we've used to encode values of the volume potential. The MT paper all but disappeared after it was submitted to and rejected by the Synposium on Theory of Computing (STOC) sometime in the early 1990s. We haven't been able to find a complete copy of this paper—all we have is a garbled fax. It's missing the first few pages, so we're not certain just what their result was, but from the email message from Thurston shown in Figure 9 it doesn't appear that they had proven the full tetvol conjecture, just the version for v large enough, as in STT.

The MT paper was an 'extended abstract'; It omitted certain details meant to be covered in the 'full paper' to follow. We haven't tried to fill in the details, but based on our own experience with this method, we have no reason to doubt that their work was essentially correct.

The difference between our approach and MT is that we deal with simple examples where we can produce an explicit volume potential. MT dealt with more complex examples, and used max-flow min-cut to find the volume potentials. We expect that their method will generalize in a way that ours will not.

5 Linear programming

Pick the volume potential ρ so as to maximize $\operatorname{vol}_{\rho}(\sigma)$, and call the maximum value $\operatorname{LPvol}(\sigma)$. STT emphasized that this is a linear programming problem:

From doyle Thu Nov 14 23:14:15 2002 To: wpt Subject: rotations

Bill,

I understand that there was a draft or a preprint related to the attached abstract. If you can send anything (e.g. a tex source) I'd love to see it.

Peter

Claire Kenyon (ENS-Lyon and William Thurston, MSRI)

Rotation distance between binary trees: hyperbolic geometry vs. max-flow min-cut.

The maximum number of rotations needed to go from one given binary tree with \$n\$ nodes to another is exactly \$2n-6\$ when \$n\$ is large enough. We first sketch Sleator, Tarjan, and Thurston's original proof of this theorem, which involves hyperbolic geometry volume arguments, the Riemann mapping theorem, approximate calculations of integrals and an induction argument. We then present an alternate, elementary proof, based on the max-flow min-cut theorem. Finally, we compare the two proofs and show how they are essentially two versions of the same proof, by relating successively hyperbolic volume to cocycles to linear programming to amortized analysis to flow problems.

From wpthurston@mac.com Fri Nov 15 01:25:13 2002
Date: Thu, 14 Nov 2002 22:23:40 -0800
Subject: Re: rotations
From: wpthurston@mac.com
To: "Peter G. Doyle" <doyle@hilbert.dartmouth.edu>

Hi Peter, Yes, there was a draft, but in all my moving around I don't have a copy any more. I should try to get my own copy from Claire. We kind of dropped it when it was rejected from STOC.

The idea was, given a triangulation that appears to be maximal, attempt to construct an L`infty 3-cocycle on \Delta^{v-1} (if there are v vertices) that takes maximal value on all the tetrahedra in your triangulation. Of course, hyperbolic volume (given an immersion of the polyhedron into H^3) gives a cocycle that works well enough in many cases --- but it's not quite the best. I don't remember all the details, although I could reconstruct them. I think it turned out that the important values to work out were when 2 or more vertices of the tetrahedron are connected by an edge; these could be done using many instances of a max-flow min-cut process: I think, one instance for each possible location for the pair of non-adjacent vertices (it was like a flow from one of these vertices to the other). When three vertices are all mutually connected by edges (i.e. the tetrahedron has a face on the bounding sphere) I think there is some formula that we just wrote down. I think the cocyle could assign 0 to many of the other tetrahedra, the ones with 4 non-adjacent vertices.

Of course a cocycle like this is really the dual to the L^1 3-chain having boundary the given triangulation. There might or might not be a geometric triangulation realizing this minimum, but it seemed to work out in lots of cases, including explicit examples for every value of the number of vertices such as in the original paper. Bill

Figure 9: The word from Bill

it's dual to the problem of minimally tetrating σ with fractional tets allowed.

The value of **LPvol** is not always an integer. (See section 6.) When it isn't, we can round up:

$$\mathbf{tetvol}(\sigma) \geq [\mathbf{LPvol}(\sigma)].$$

We have yet to find a triation where this doesn't give us **tetvol**:

Hypothesis (LPvol).

$$\mathbf{tetvol}(\sigma) = [\mathbf{LPvol}(\sigma)].$$

This is meant as a null hypothesis, not a conjecture. The evidence we've accumulated so far is extensive but rather weak. ('It seemed to work in a lot of cases.') The problem is that things only start to get interesting when you get to v = 20 or so, and we have difficulty computing **tetvol** or even **LPvol** for v much bigger than this.

If this hypothesis is true, computing **tetvol** would reduce to linear programming, and would thus take (weakly) polynomial time.

6 More about phyllohedra

Let's quickly indicate what happens when we truncate other phyllocylinders, keeping the details for another day.

The phyllocylinder $\Phi_{a,b}$ has combinatorial girth a + b. Along with $\Phi_{5,1}$, the other phyllocylinders of girth 6 are $\Phi_{6,0}$, $\Phi_{4,2}$, and $\Phi_{3,3}$. $\Phi_{6,0}$ is an infinite stack of hexagonal antiprisms. Like $\Phi_{5,1}$, $\Phi_{6,0}$ has volume-to-surface-area ratio 1, meaning that it has a volume potential ρ taking value 1 to each of its triangles: $\mathbf{vsa}(\Phi_{6,0}) = 1$. Figure 10 shows the associated flow. (As usual it falls short in that it doesn't indicate values for triangles not containing an edge.)

By contrast,

$$\mathbf{vsa}(\Phi_{4,2}) = \frac{31}{32}.$$

Long (4, 2)-phyllohedra have **tetvol** asymptotically equal to $\frac{31}{16}v$. Truncating and capping in the most natural way (see Figure 11), we get a family U_v (v even) for which **LPvol** is not always an integer — but we still have $[\mathbf{LPvol}] = \mathbf{tetvol}$. In fact there is a volume potential ρ on the infinite



Figure 10: Volume potential flow for $\Phi_{6,0}$. Flow rate is 1 along thick arrows, 1/2 along thin arrows.



Figure 11: Truncating $\Phi_{4,2}$ and capping in the most natural way yields a family of phyllohedra U_v , v even, answering the question, 'What would the icosahedron look like if it had 14 vertices?'

v	$\mathbf{tetvol}(U_v)$	tetvol - LPvol	$tetvol - \rho$
12	15	0	7/8
14	18	0	0
16	22	0	1/8
18	26	0	1/4
20	30	1/5	3/8
22	34	1/3	1/2
24	38	1/2	5/8
26	42	2/3	3/4
28	46	2/3	7/8
30	49	0	0
32	53	0	1/8

Figure 12: A single volume potential ρ defined on $\Phi_{4,2}$ produces sharp lower bounds $\lceil \rho(U_v) \rceil$ for $\mathbf{tetvol}(U_v)$. The table shows the shortfall of the lower bounds $\mathbf{LPvol}(U_v)$ and $\rho(U_v)$. As these are less than 1, rounding up gives the exact value.

cylinder which when restricted may come in lower than the actual value of **LPvol**, but still yields **tetvol** when rounded up, giving

$$\mathbf{tetvol}(U_v) = \lceil \rho(U_v) \rceil = \lceil 31/32(2v - 12) + 5/2 \rceil.$$

(See Figure 12.)

For $\Phi_{3,3}$ the results are similar, only now with $\mathbf{vsa} = \frac{23}{24}$.

For girth 7 all **vsa** are 1. From $\Phi_{5,2}$ or $\Phi_{4,3}$ we get two new families with **tetvol** = 2v - 10. For $\Phi_{6,1}$, when you cap in the natural way there are vertices of degree 7, and **tetvol** = 2v - 11; other methods of capping give triations with **tetvol** = 2v - 10.

7 Flip distance

The work of STT [4] on tetrations was motivated by the problem of finding the maximum possible flip distance between two triangulations of a v-gon. If α, β are triangulations of a v-gon with no common edge we get a triation $\alpha - \beta$ of the 2-sphere by gluing along their common v-gon. A flip path from α to β begets a tetration of $\alpha - \beta$ so

$$\mathbf{flipdist}(\alpha,\beta) \geq \mathbf{tetvol}(\alpha-\beta)$$

Pournin [2]. gave examples with **flipdist** = 2v - 10, $v \ge 13$, thus proving the analog of the tetvol conjecture for **flipdist**. We can identify his examples as arising from truncations of $\Phi_{5,2}$, outfitted with a particular natural Hamiltonian cycle. By varying the Hamiltonian cycles we get many other examples from these same triations.

Pournin's **flipdist** result doesn't imply the tetvol conjecture for triations, because there may be a gap between **tetvol** and **flipdist**: The Hamiltonian cycle can prevent us from converting an optimal tetration into a flip path. Wang [5] gives examples where the ratio of **flipdist** to **tetvol** is arbitrarily close to 3/2. In Wang's examples **flipdist** and **tetvol** are on the order of 3/2vand v, so far from the kind of triations we're dealing with here. But we can take Wang's smallest example, which has v = 10, **flipdist** = 2v - 10 = 10, **tetvol** = 2v - 11 = 9 (Figure 13), and inflate it to an example with v = 16, **tetvol**($\alpha - \beta$) = 2v - 11 = 21, and **flipdist**(α, β) = 2v - 10 = 22. (Figures 14,15.) This illustrates why knowing examples with **flipdist** = 2v - 10doesn't settle the tetvol conjecture.



Figure 13: The 10-vertex triation on the left has tetvol = 2v - 11 = 9. The black Hamiltonian cycle divides it into disk triangulations α (blue) and β (red), with flipdist(α, β) = 2v - 10 = 10.



Figure 14: A larger example. This time v = 16, $tetvol(\alpha - \beta) = 2v - 11 = 21$, and $flipdist(\alpha, \beta) = 2v - 10 = 22$.



Figure 15: The v = 16 example brought to life.

8 The 3-ball

We've sidestepped the question of whether a minimal tetration of a triation σ of the 2-sphere yields a triangulation of the 3-ball. It does. That's because there are fewer tets than faces, so some tet must meet σ in at least two faces, necessarily adjacent. Remove this tet, and you get a minimal triation of a sphere, or a pair of spheres (or nothing at all, if you were down to a single tet). Proceed by induction.

Contrast this with what happens for triations of a torus, or a surface of higher genus. In that case there is no guarantee that a minimal tetration will be a manifold, or even a pseudo-manifold. And if it is a manifold, we have no a priori control over how it fills in the surface. All very mysterious.

9 What's true in general

If a triation has all vertices of degree 5 or 6, the chances are that $\mathbf{tetvol} = 2v - 10$. The only exceptions we know are phyllohedra derived from $\Phi_{4,2}$ and $\Phi_{3,3}$. (This includes the icosahedron, which can be derived from either.) We'll stop short of formulating a precise conjecture. The point is that producing 2v - 10 triations is not the issue, it's proving that they have this property. As stated above, we expect that the right approach is that of MT.

References

- Claire Mathieu Kenyon and William P. Thurston. Rotation distance between binary trees: hyperbolic geometry versus max-flow min-cut. (Title is approximate), 1992 or 1993.
- [2] Lionel Pournin. The diameter of associahedra. Adv. Math., 259:13-42, 2014. https://doi.org/10.48550/arXiv.1207.6296.
- [3] Daniel D. Sleator, Robert E. Tarjan, and William P. Thurston. Rotation Distance, Triangulations, and Hyperbolic Geometry. In *Proceedings of the* 18th Annual ACM Symposium on Theory of Computing, pages 122–135. ACM, 1986.

- [4] Daniel D. Sleator, Robert E. Tarjan, and William P. Thurston. Rotation distance, triangulations, and hyperbolic geometry. J. Amer. Math. Soc., 1(3):647–681, 1988.
- [5] Zili Wang. Flip distance and triangulations of a ball, 2022, arXiv:2205.11623 [math.GT]. https://doi.org/10.48550/arXiv. 2205.11623.