Normal Surfaces

Definition. Take a manifold *M* with a fixed triangulation.
A *normal surface* is a surface that intersects each tetrahedron in a finite collection of disjoint triangles and quadrilaterals

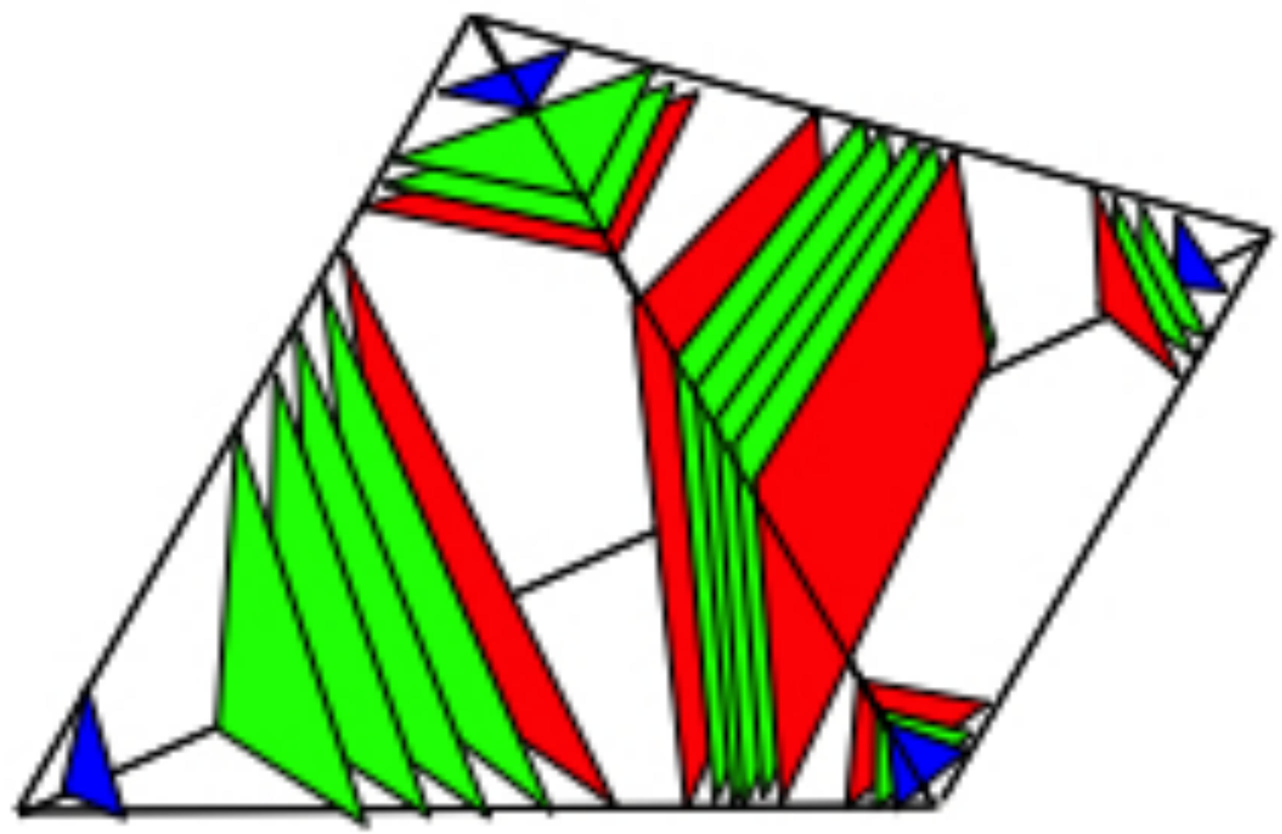

Normal Surfaces

Definition. A normal surface is a surface that intersects each tetrahedron in a finite collection of triangles and quadrilaterals.

Each tetrahedron admits four types of triangle and three types of quadrilateral.

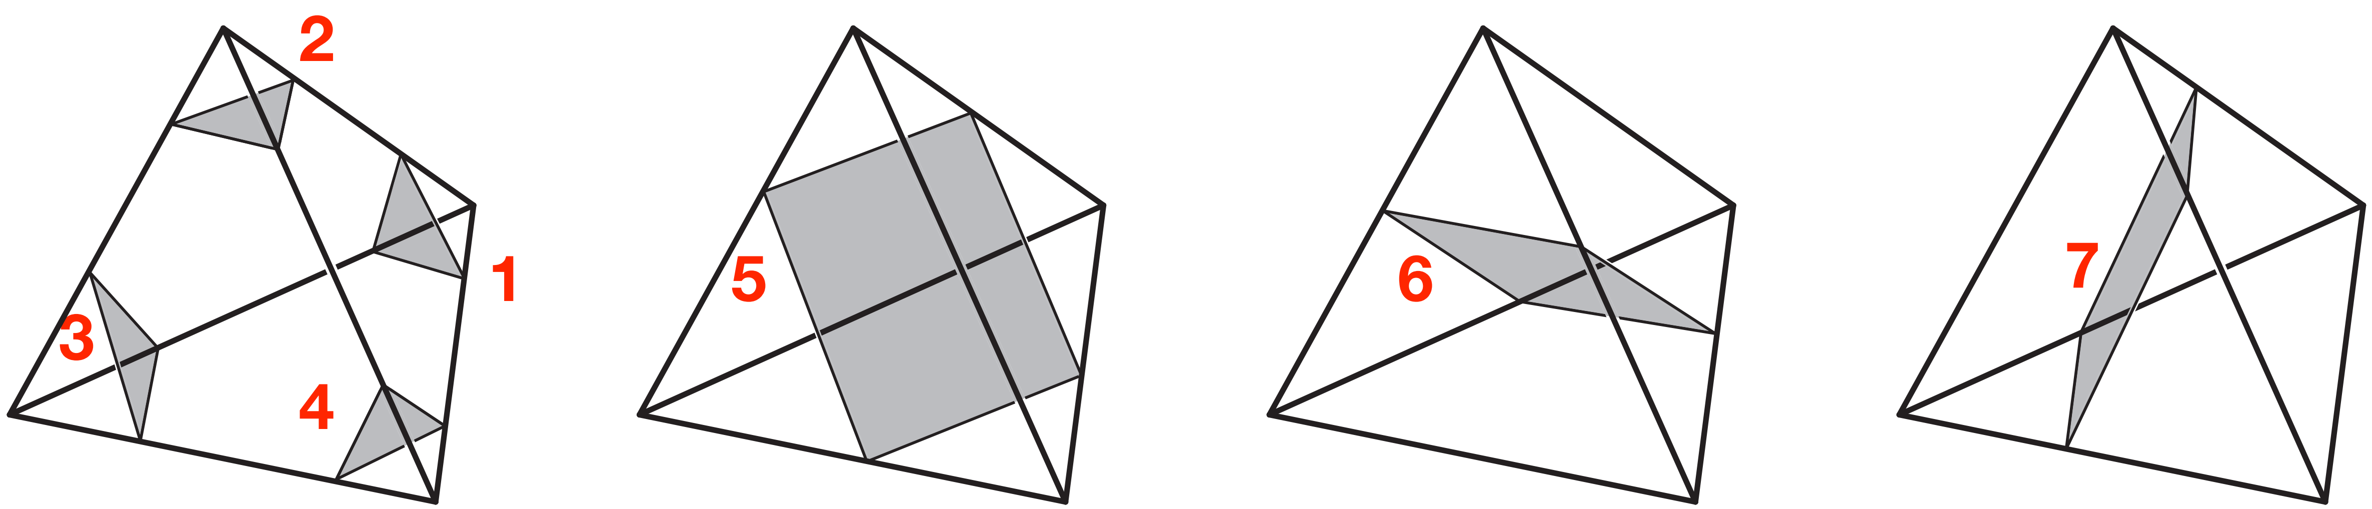

Normal Surfaces

Definition. A normal surface is a surface that intersects each tetrahedron in a finite collection of triangles and quadrilaterals.

Each tetrahedron admits four types of triangle and three types of quadrilateral.

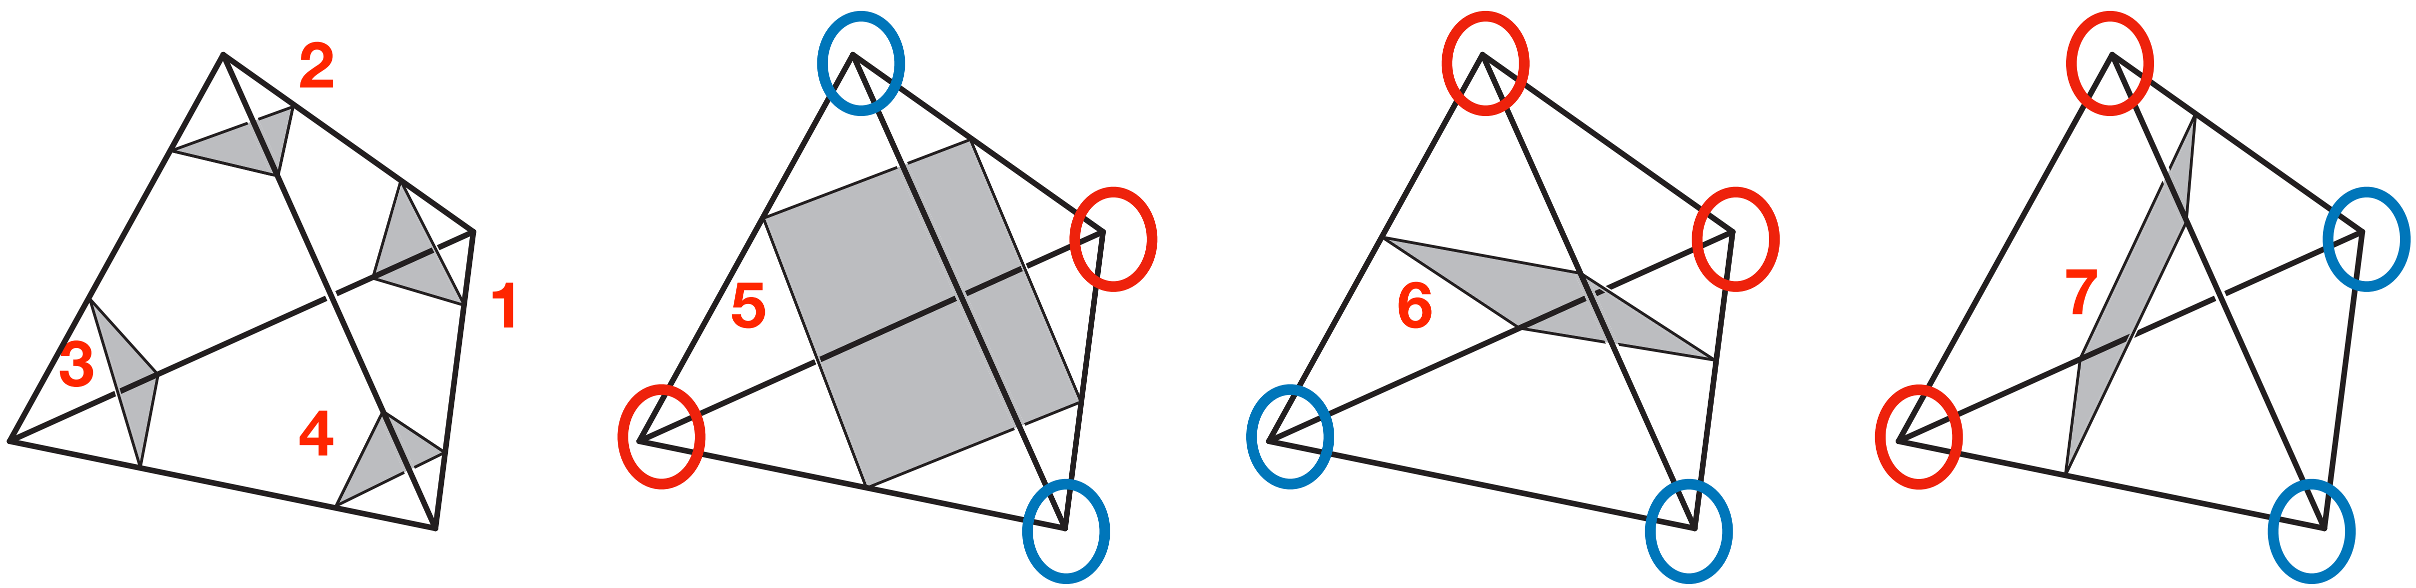

Distinct quadrilateral types separate different pairs of vertices.
Distinct types cannot be made disjoint. They must intersect.

An embedded surface can only have one type of quadrilateral in each tetrahedron.

Normal Surfaces

Definition. A normal surface is a surface that intersects each tetrahedron in a finite collection of disjoint triangles and quadrilaterals.

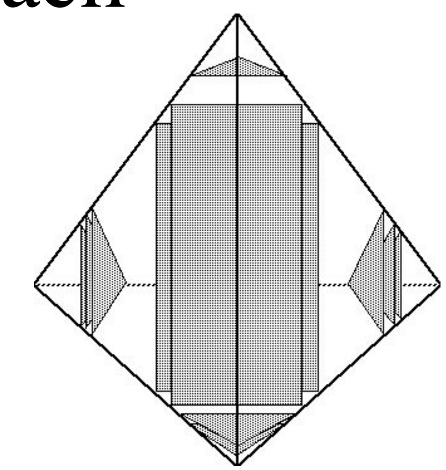

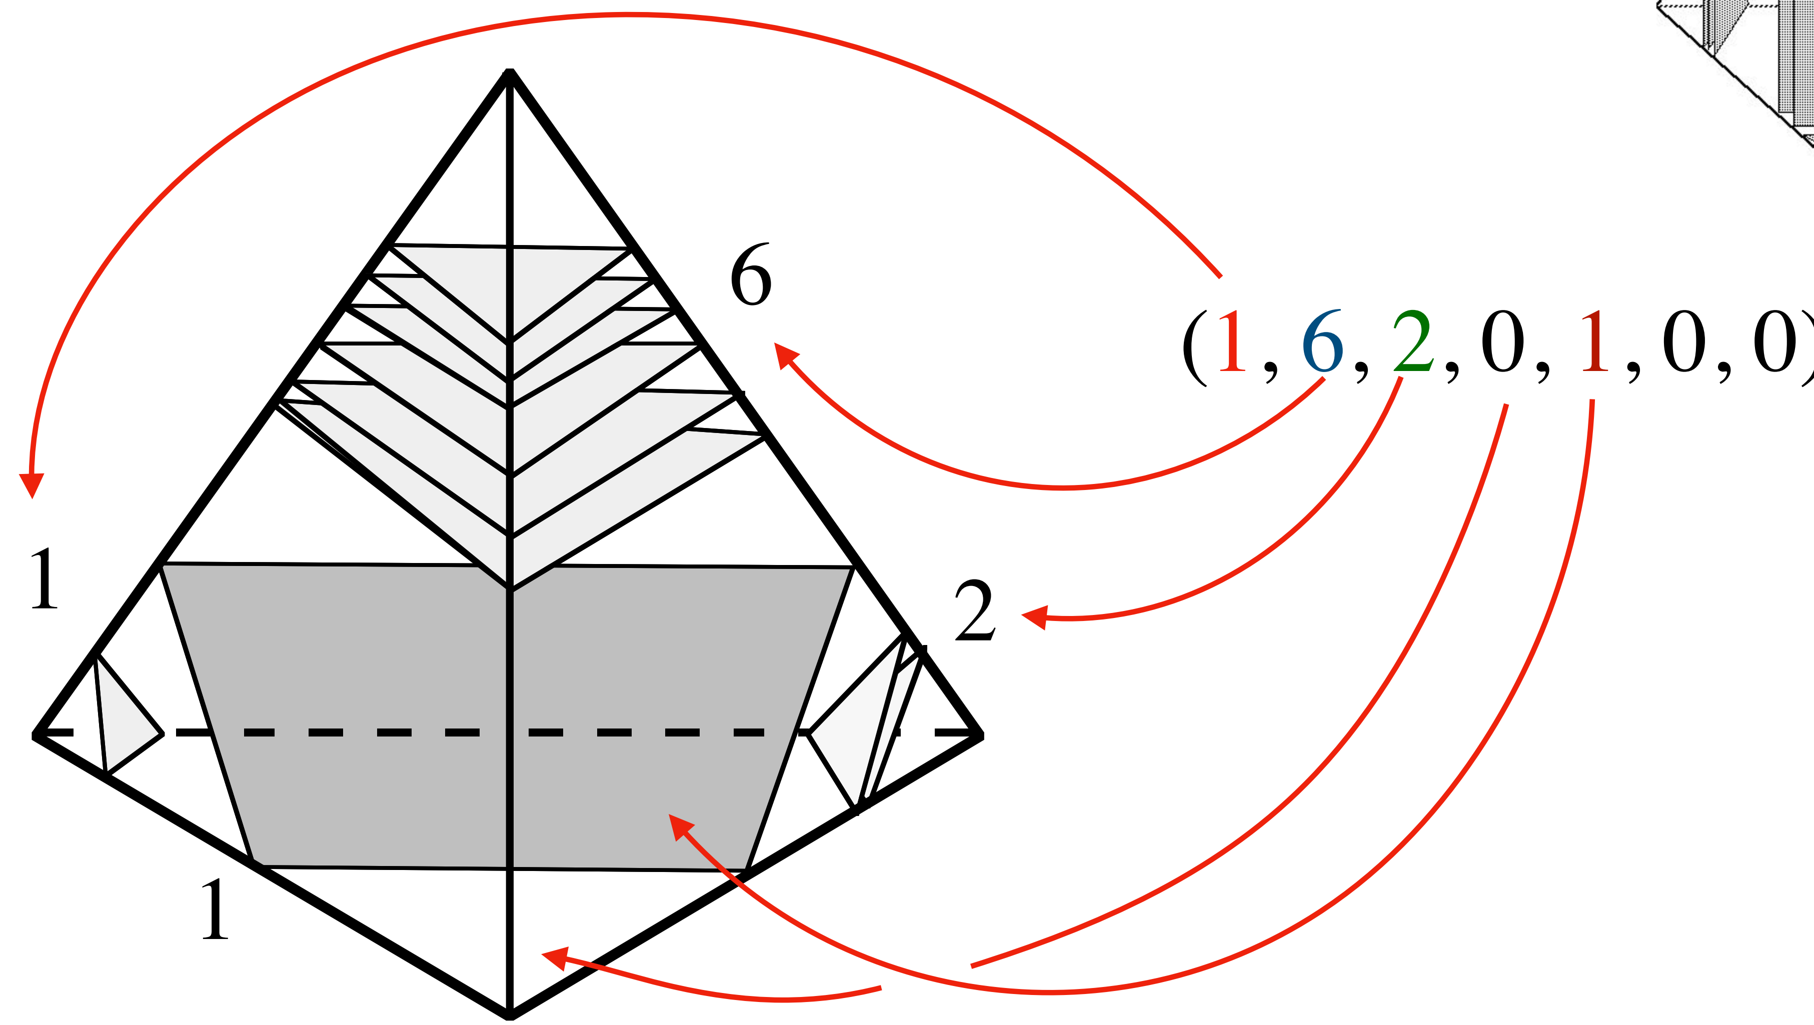

There can be many parallel copies of each type of triangle and quadrilateral. These are counted by a vector with seven coordinates.

Quadrilateral types

Distinct quadrilateral types *must* intersect. So in the correspondence between embedded surfaces and vectors, we look at vectors for which at least two of the three quadrilateral coordinates in each tetrahedron are zero

This causes a bookkeeping issue. To avoid dealing with it, we consider 3^t special cases of our problem. In each tetrahedron, we look at normal surfaces that have one choice of quadrilateral type specified.
This multiplies the complexity of our problem by 3^t, an exponential increase.
All other aspects of problems including UNKNOTTING and SPLIT LINK are polynomial (Casson). This step is the only obstacle in getting Haken's algorithm to polynomial running time.

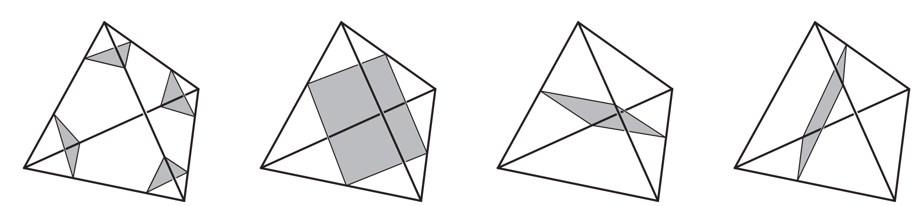

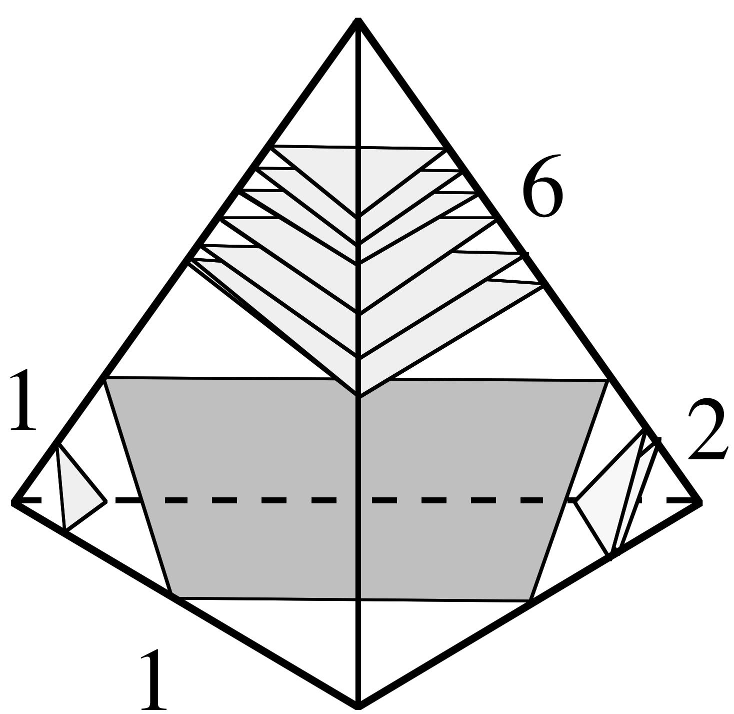

(1, 6, 2, 0, 1, 0, 0)

Set these two quadrilateral coordinates to be zero.

Normalization

Many classes of surfaces in a triangulated 3-manifold can be deformed until they are normal.

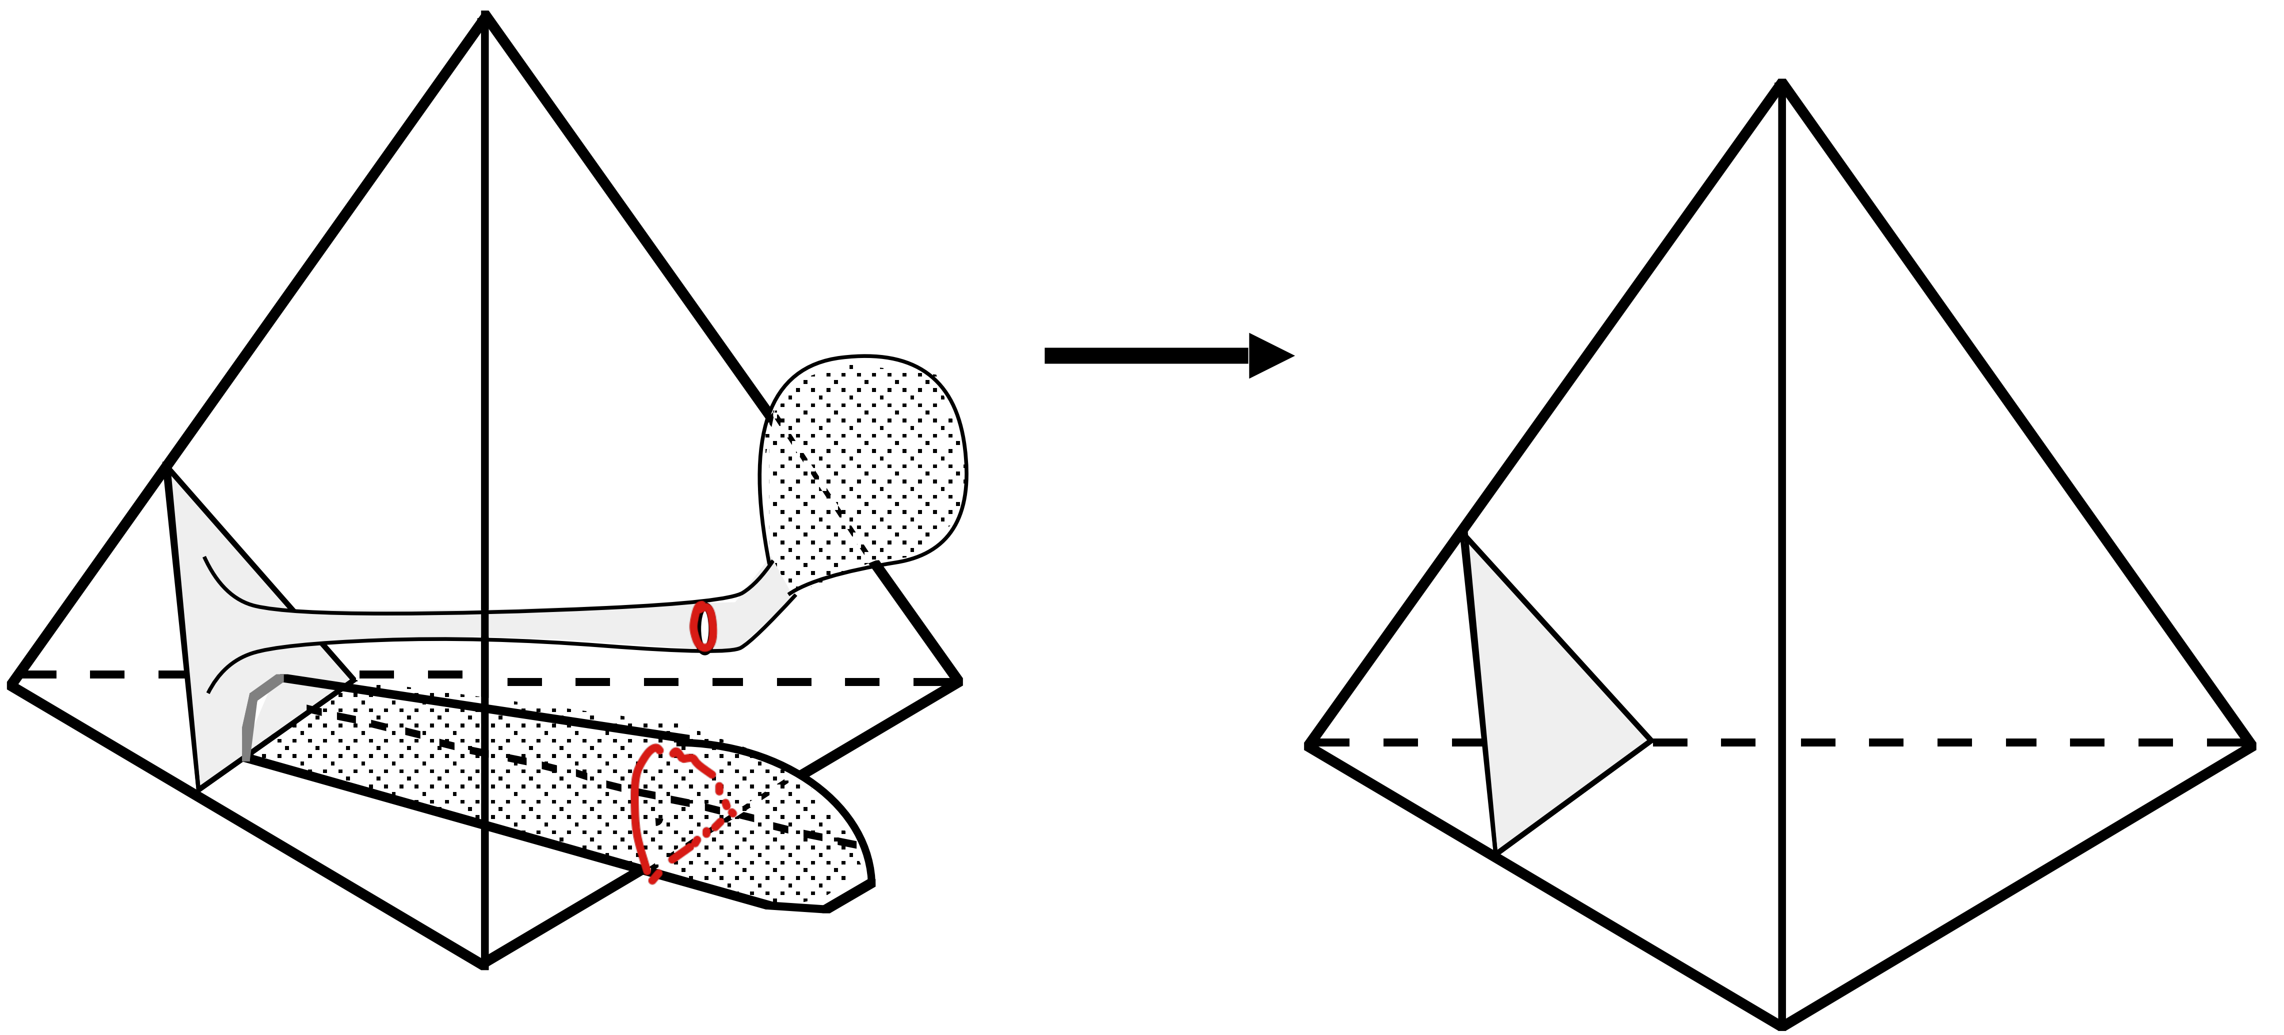

Which classes?
We first look at how to turn an arbitrary surface into a normal surface.

Normalization

Theorem Any surface F in a triangulated 3-manifold M can be transformed to a normal surface by a sequence of the following moves:

1. Isotopy
2. Compression and boundary-compression
3. Eliminating components lying inside a single tetrahedron.

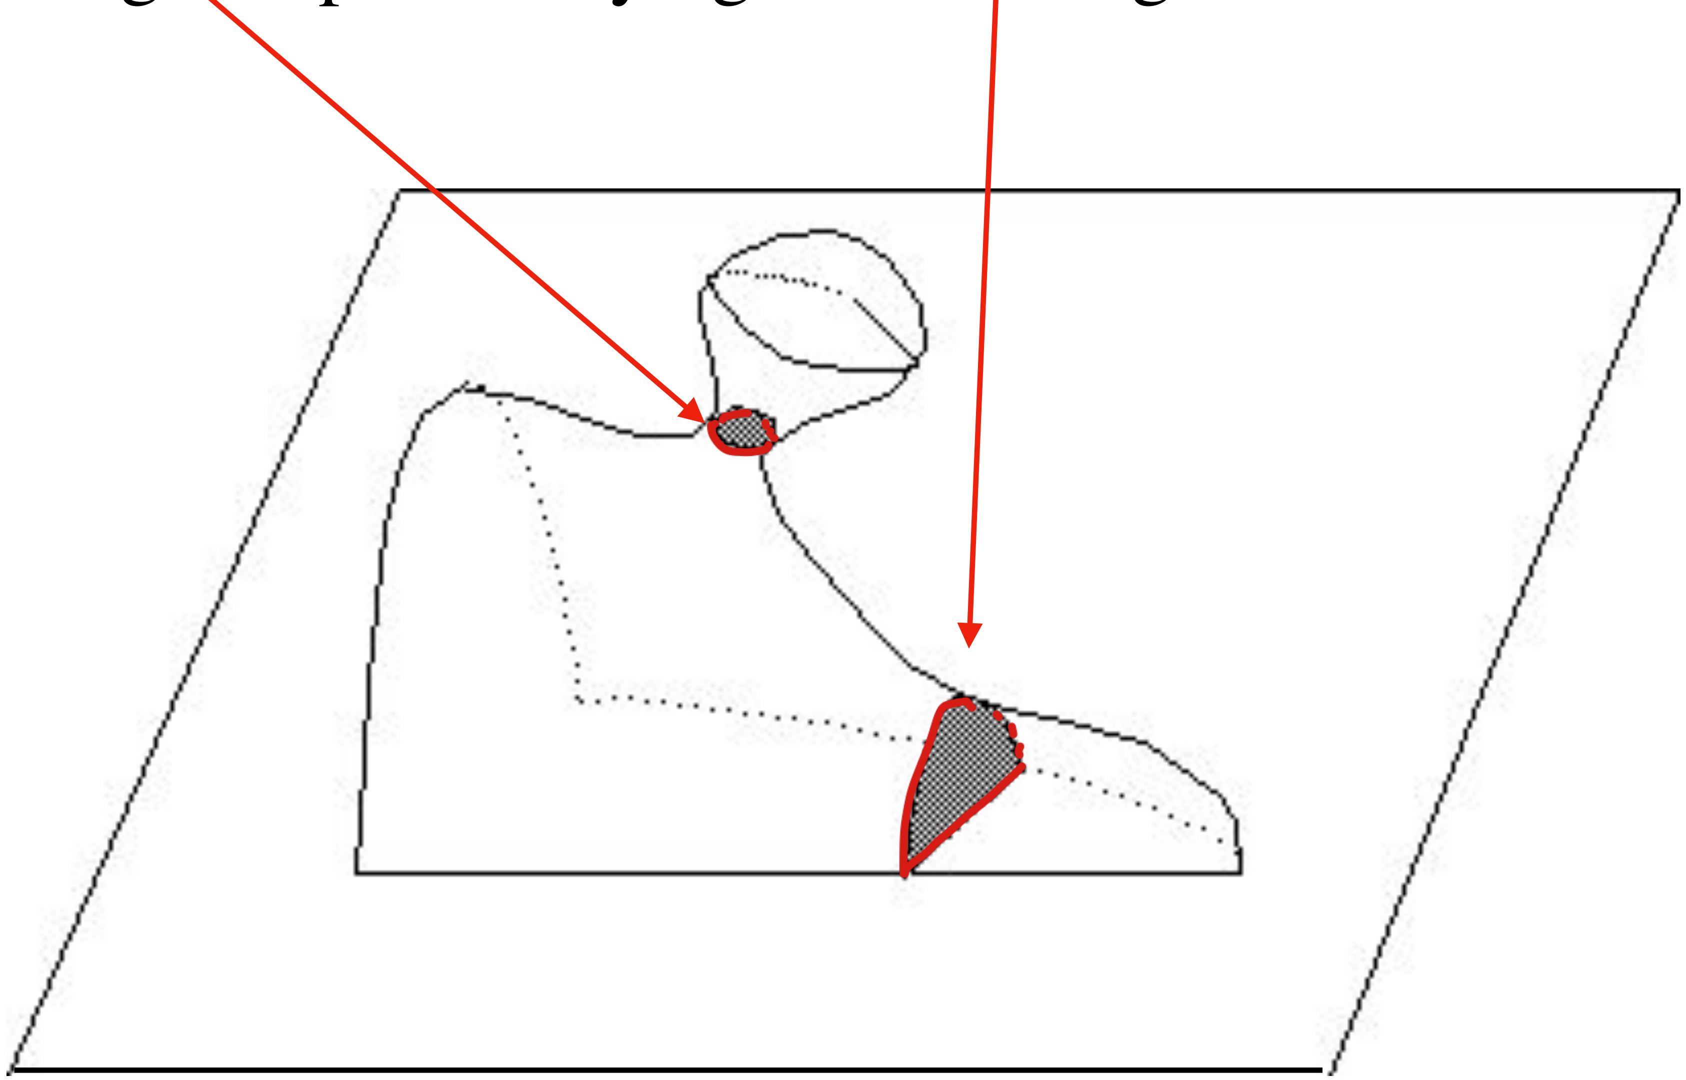

Normalization

Theorem Any surface F in a triangulated 3-manifold M can be transformed to a normal surface by a sequence of the following moves:

1. Isotopy
2. Compression and boundary-compression
3. Eliminating components lying inside a single tetrahedron.

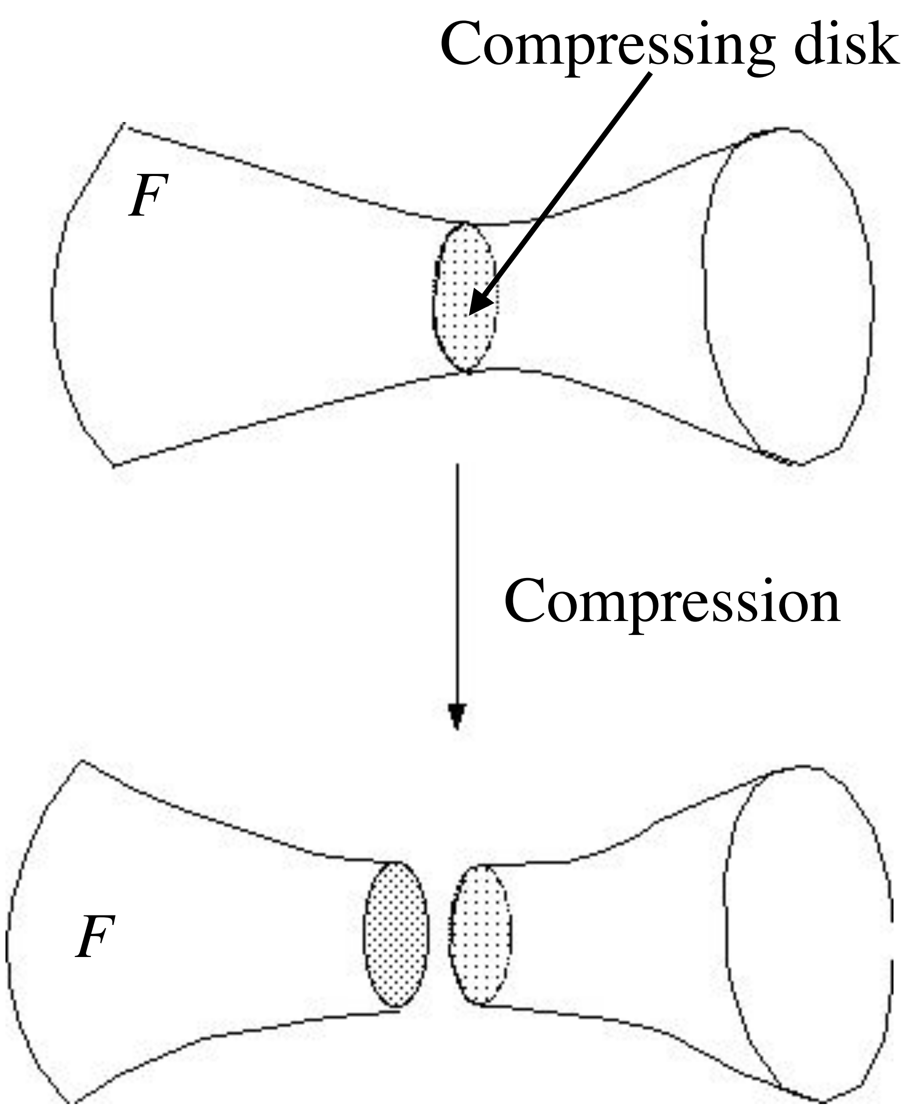

Compression of a surface $F \subset M$

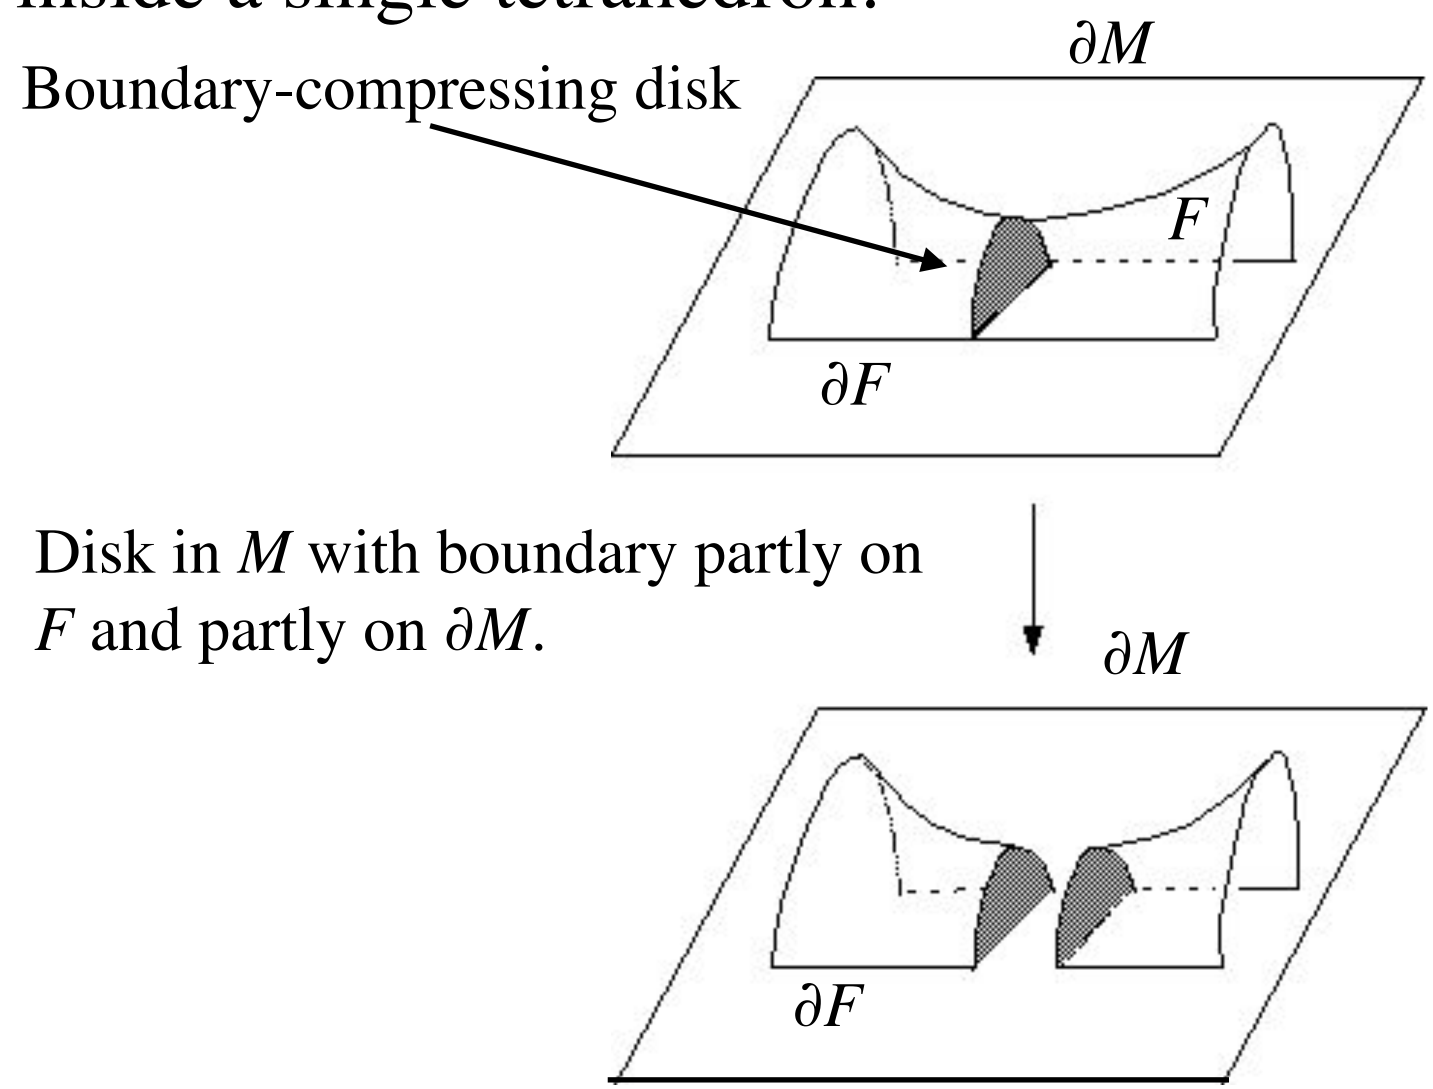

Boundary-compression of a surface $(F, \partial F) \subset (M, \partial M)$

Normalization

Theorem Any surface F in a triangulated 3-manifold M can be transformed to a normal surface by a sequence of the following moves:

1. Isotopy
2. Compression and boundary-compression
3. Eliminating components lying inside a single tetrahedron.

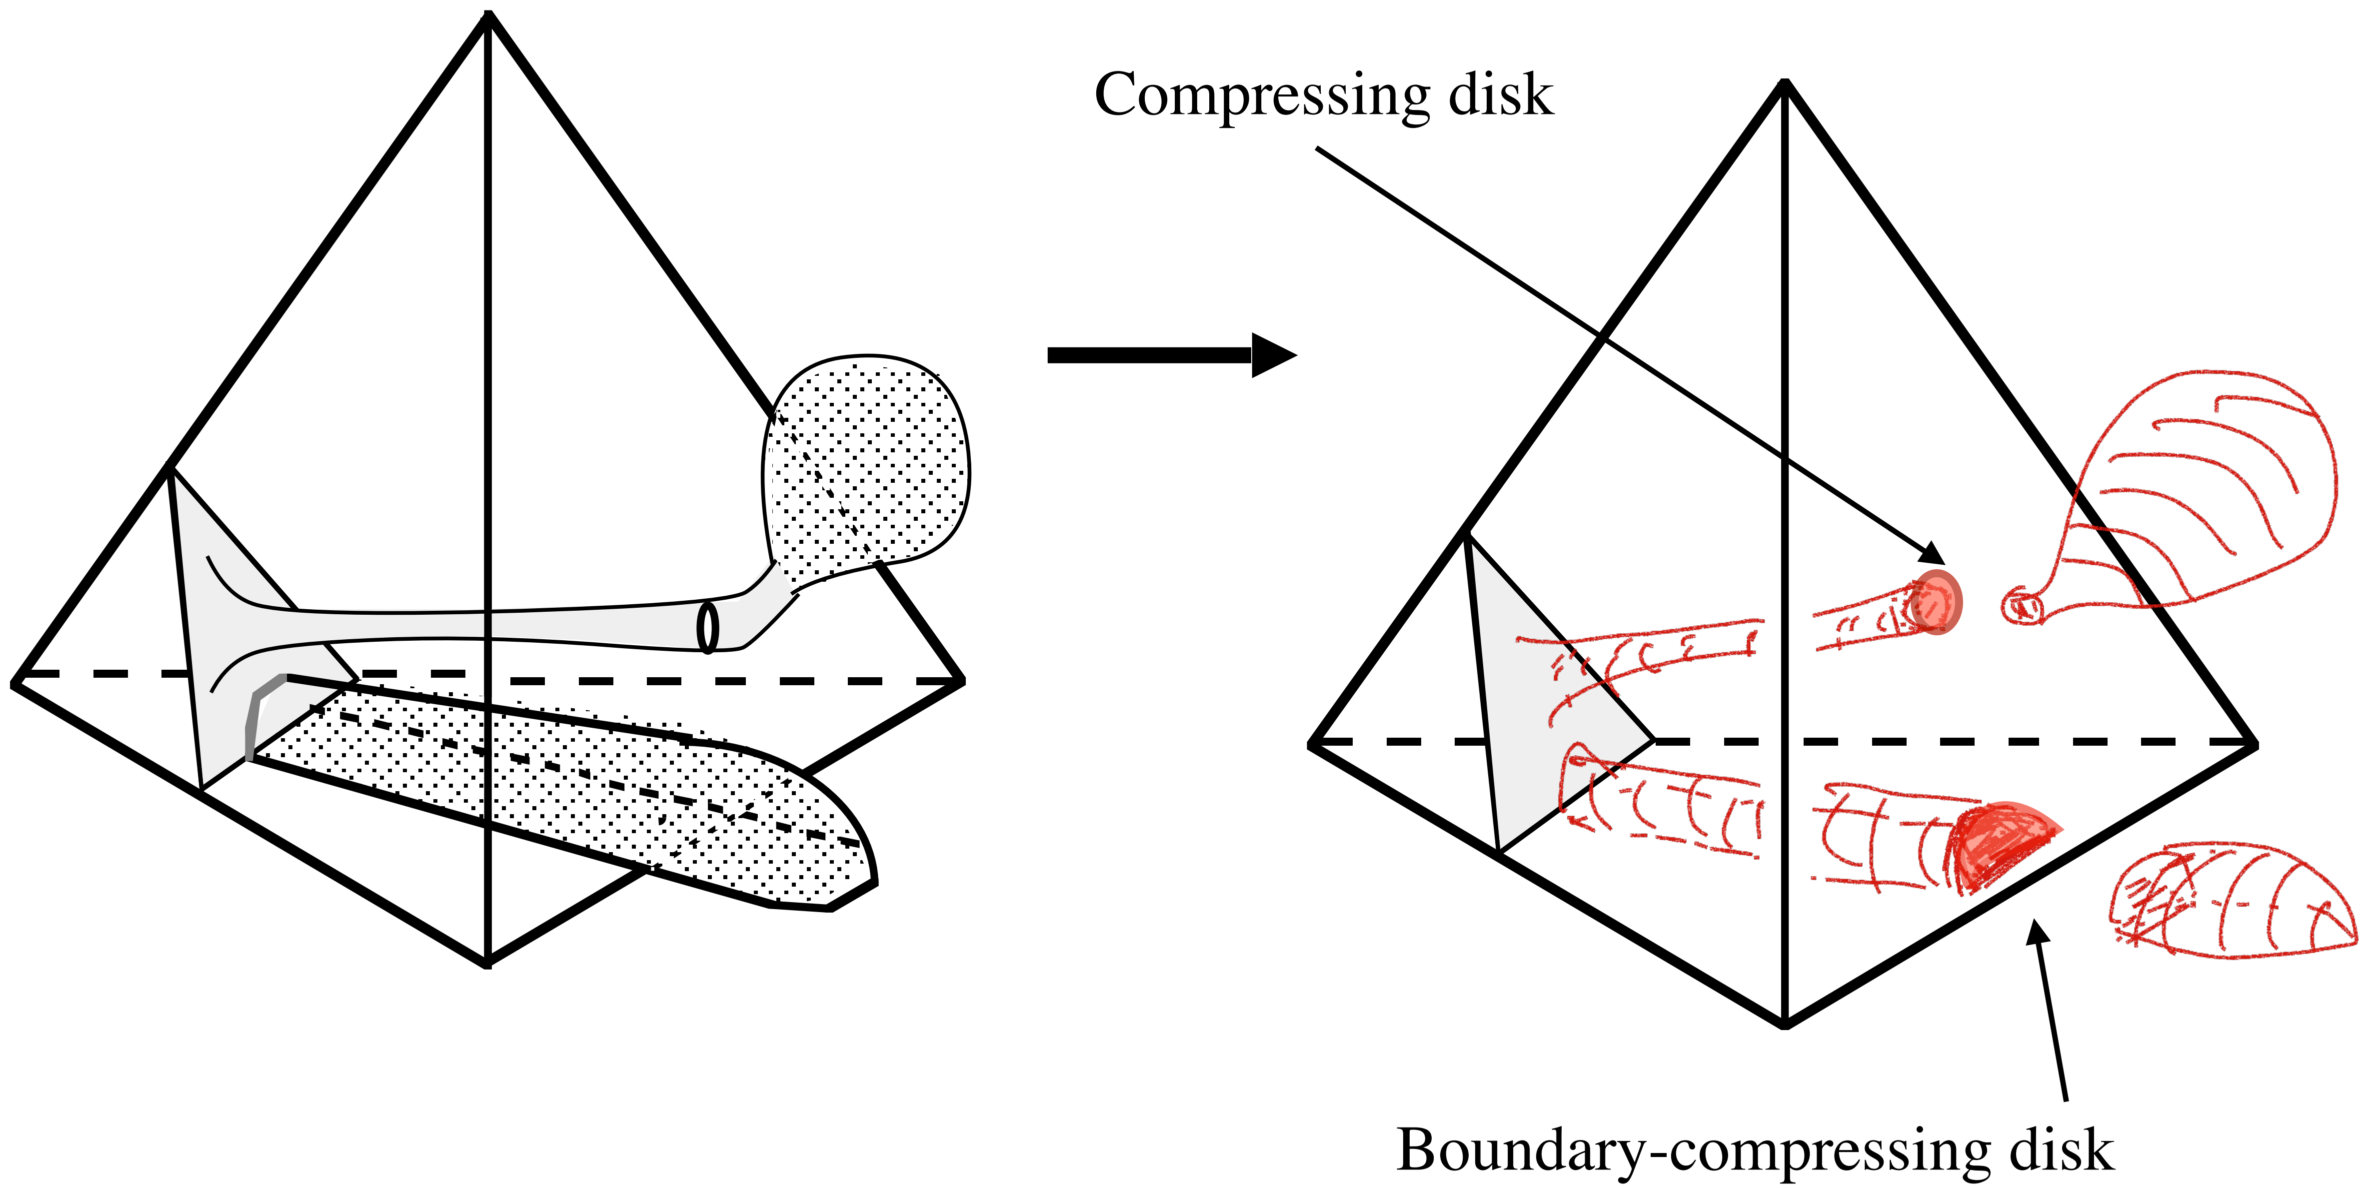

Normalization

Theorem Any surface *F* in a triangulated 3-manifold *M* can be transformed to a normal surface by a sequence of the following moves:

1. Isotopy
2. Compression and boundary-compression
3. Eliminating components lying inside a single tetrahedron.

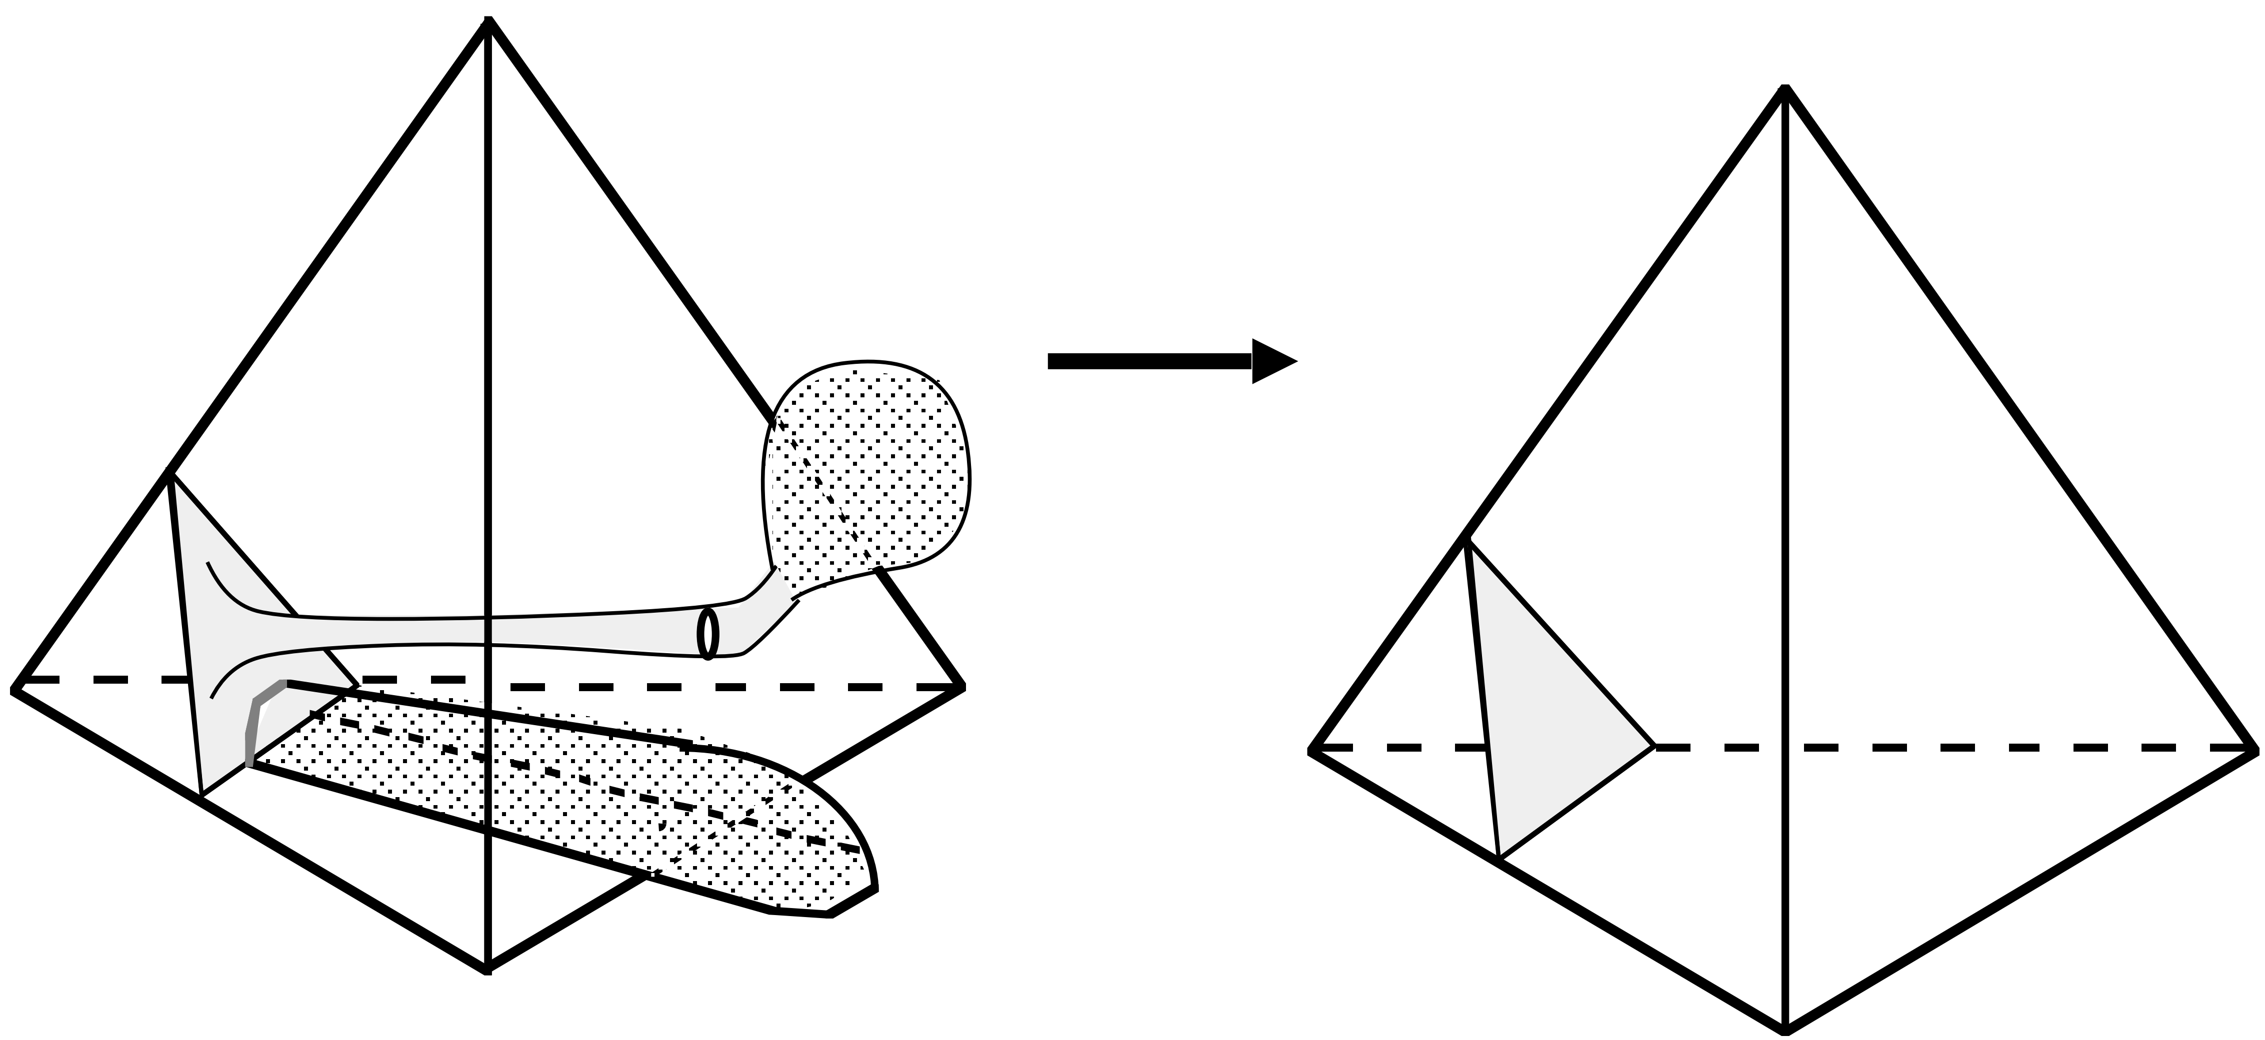

A series of such simplification moves yields a normal surface.

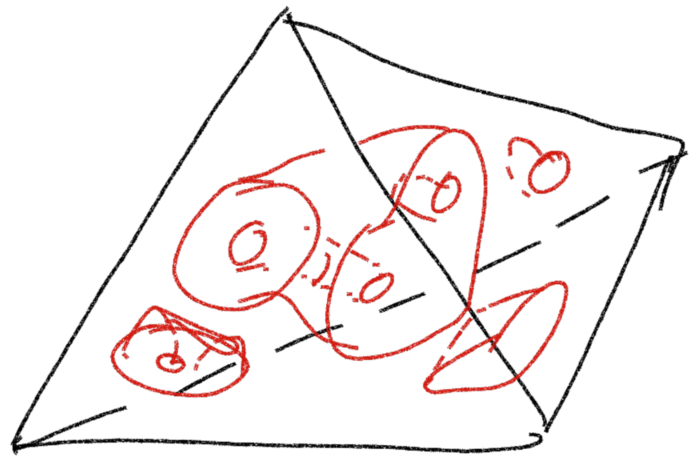

A surface can intersect a tetrahedron in a complicated way.

1. Eliminate intersection curves contained in a face of a tetrahedron.

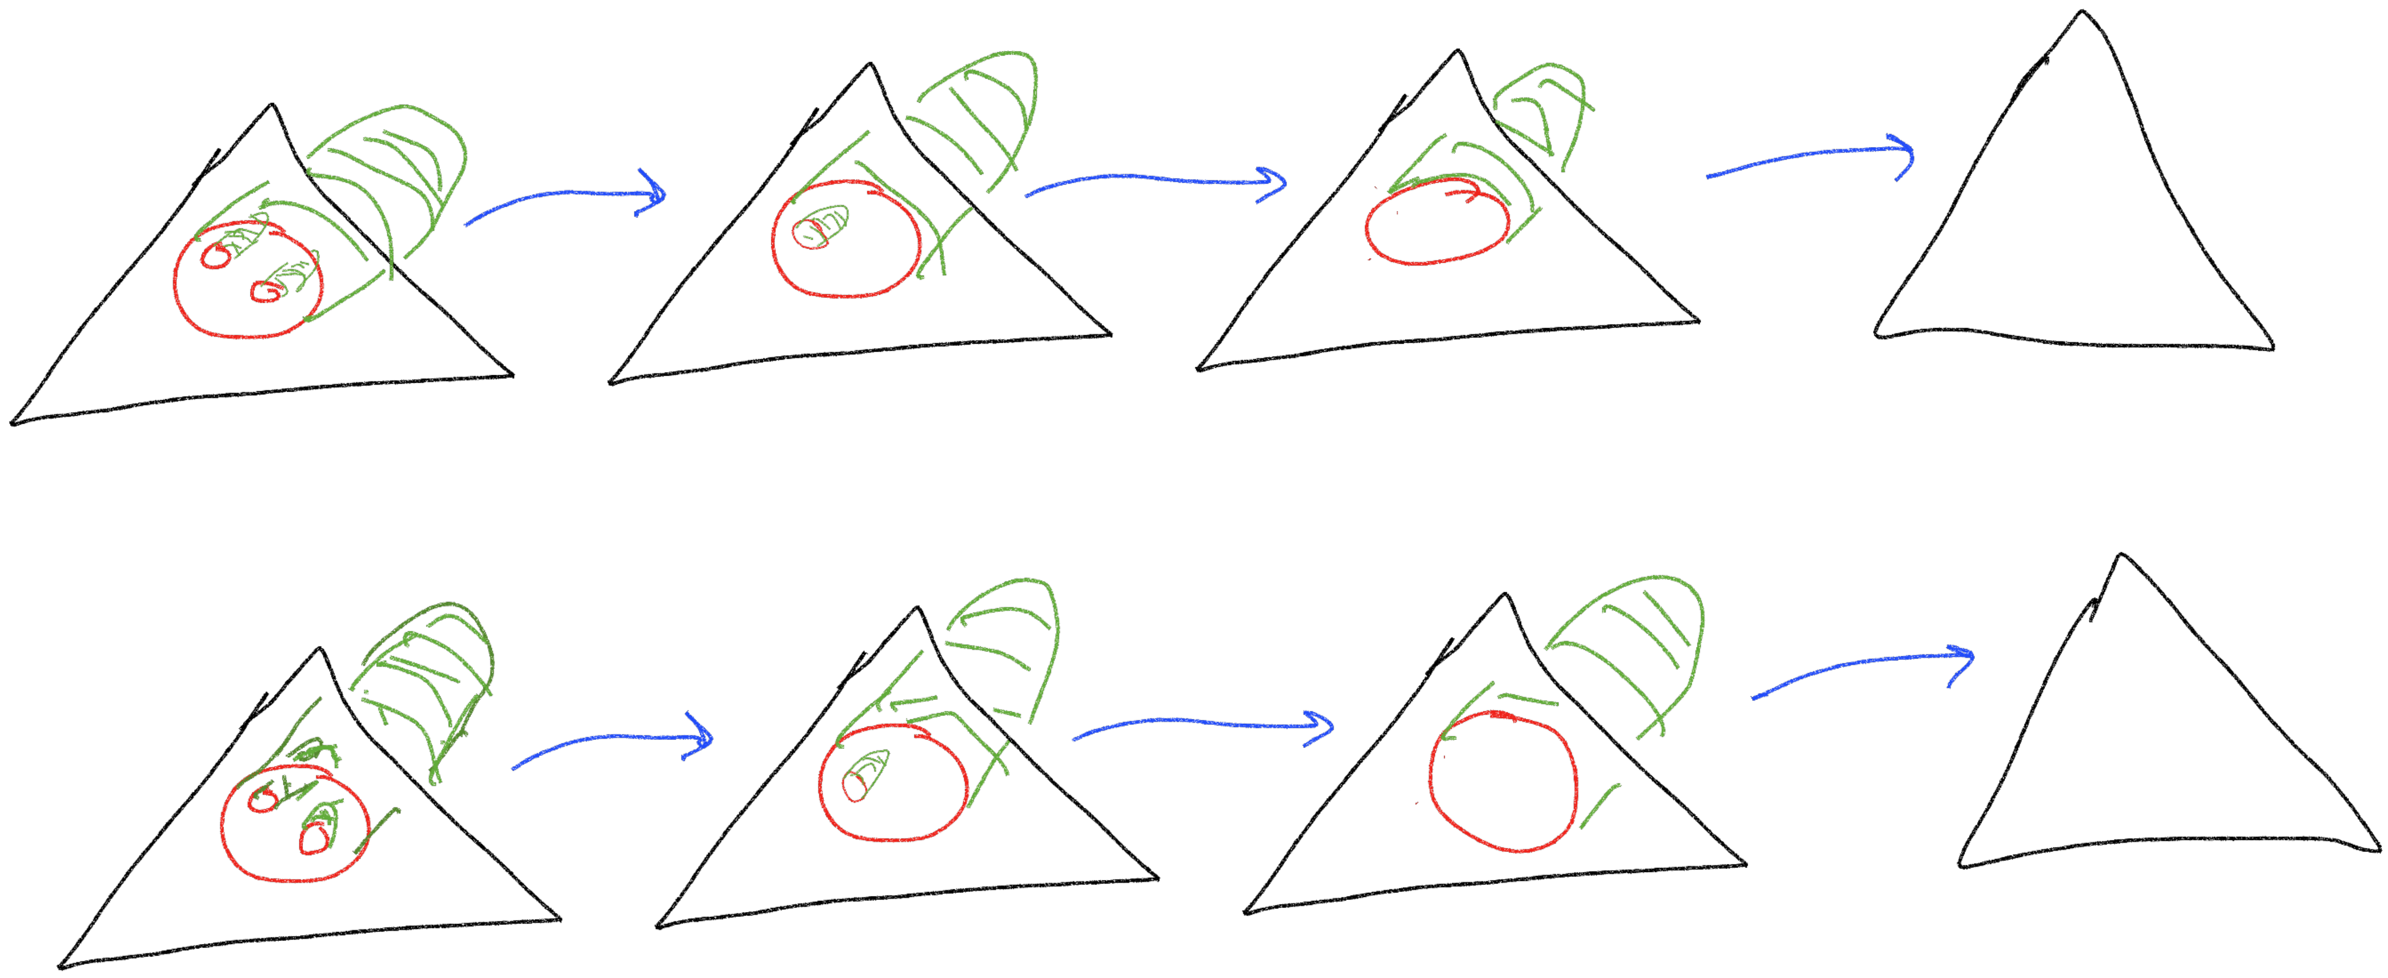

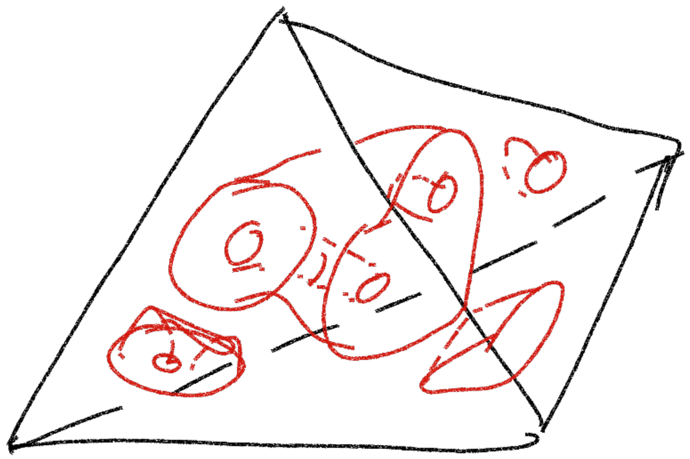

A surface can intersect a tetrahedron in a complicated way.

1. Eliminate intersection curves contained in a face of a tetrahedron.
2. Eliminate curves that meet an edge of a tetrahedron once.

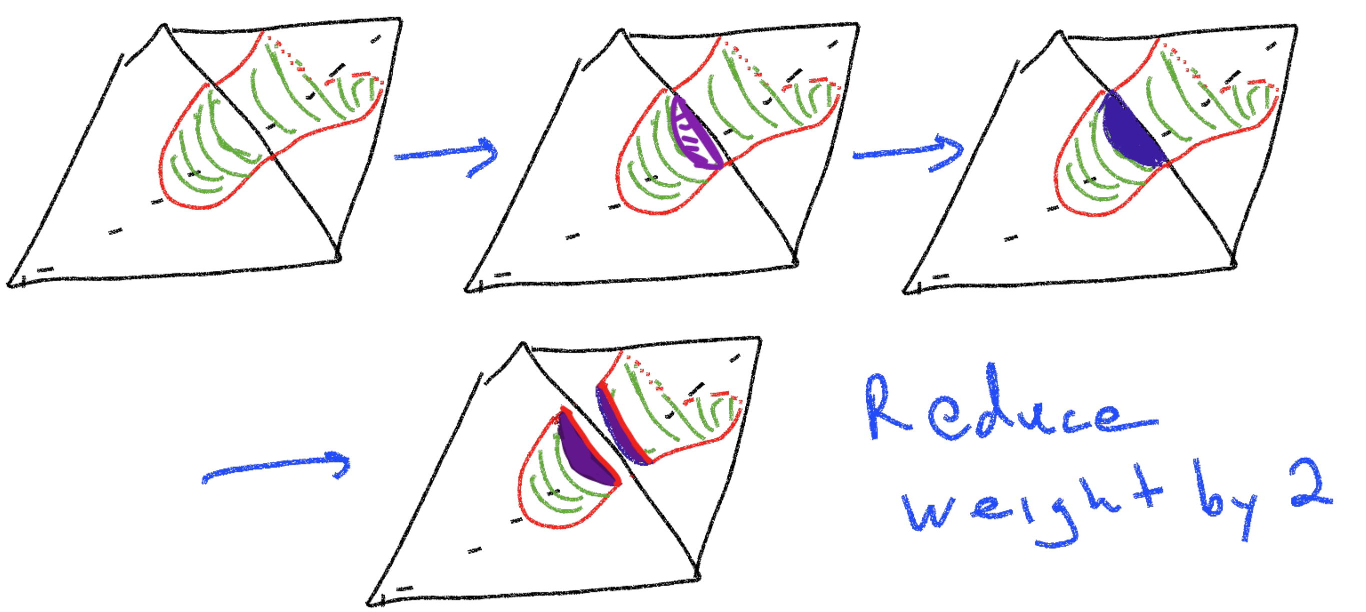

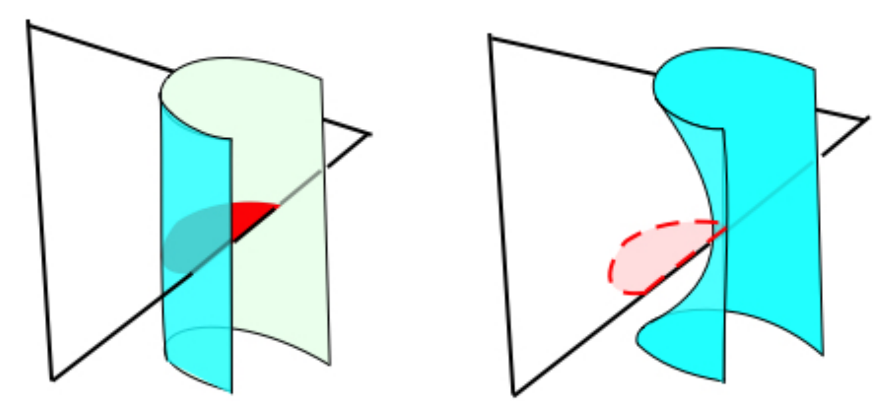

Normalization

Proof. Perform a series of moves on F that

1. Eliminate curves of intersection with a tetrahedron face by compression.
2. Eliminate arcs of intersection with a tetrahedron face by a boundary compression.

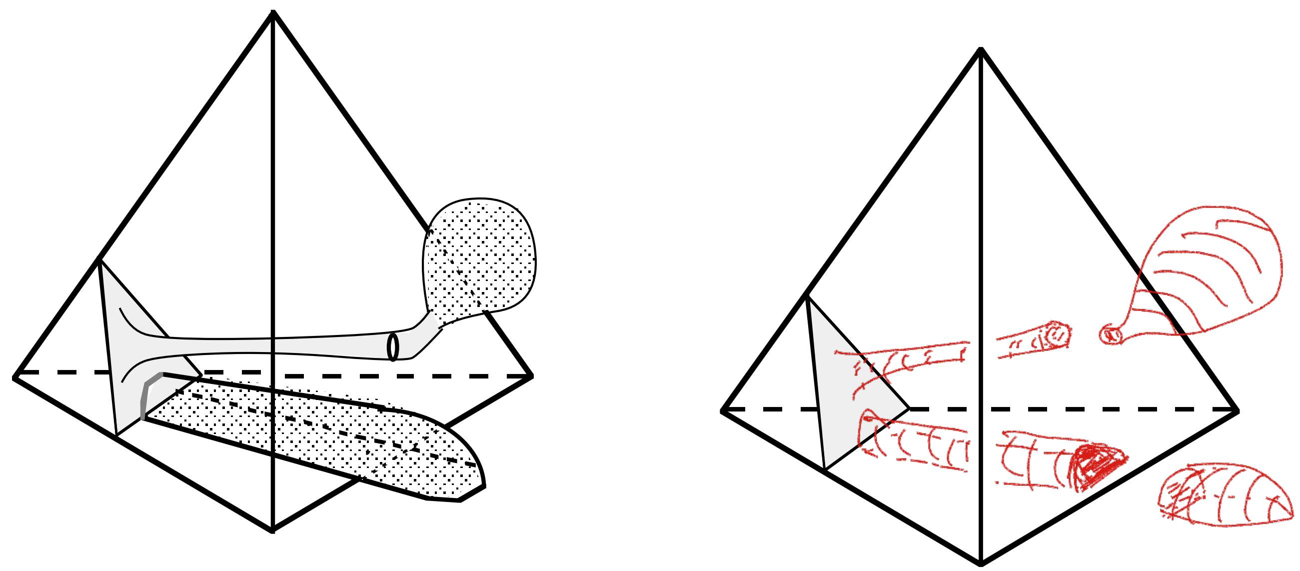

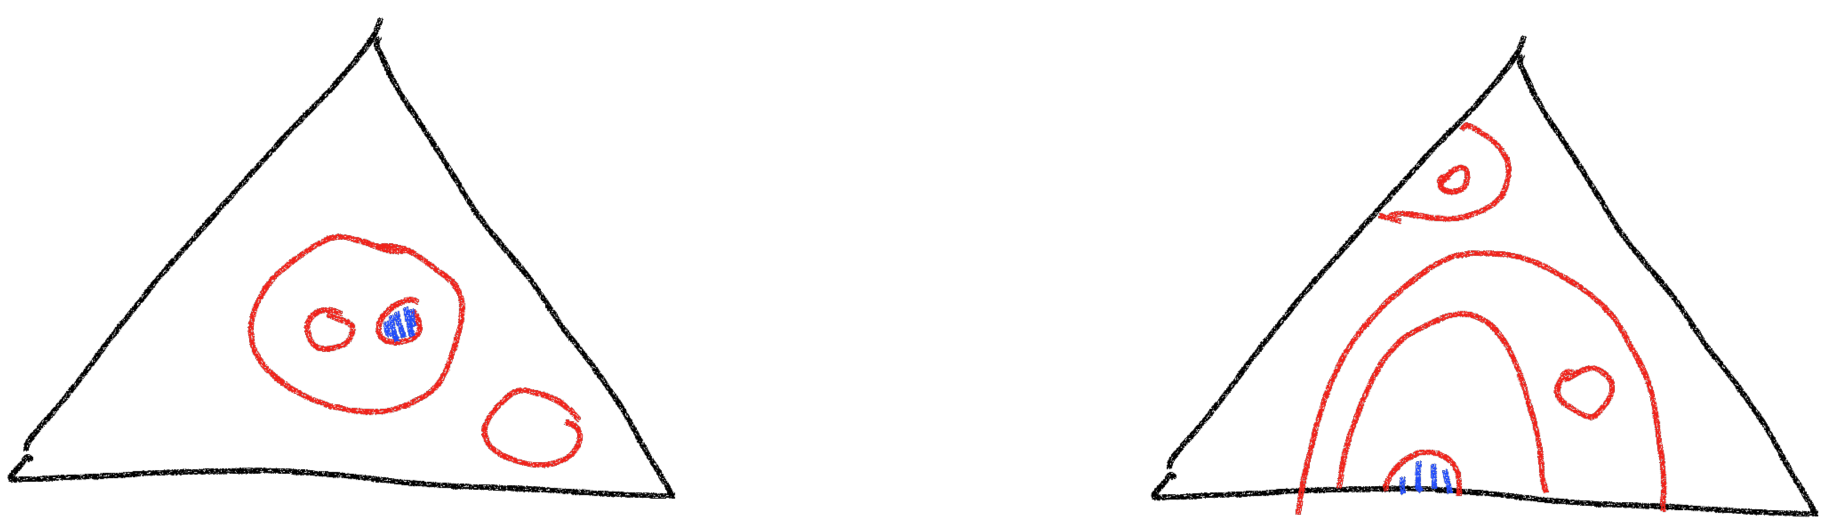

3. Not done

More Normalization

Isotop *F* to eliminate pairs of intersection points with a single tetrahedron edge. (If there are several, eliminate an innermost pair first)

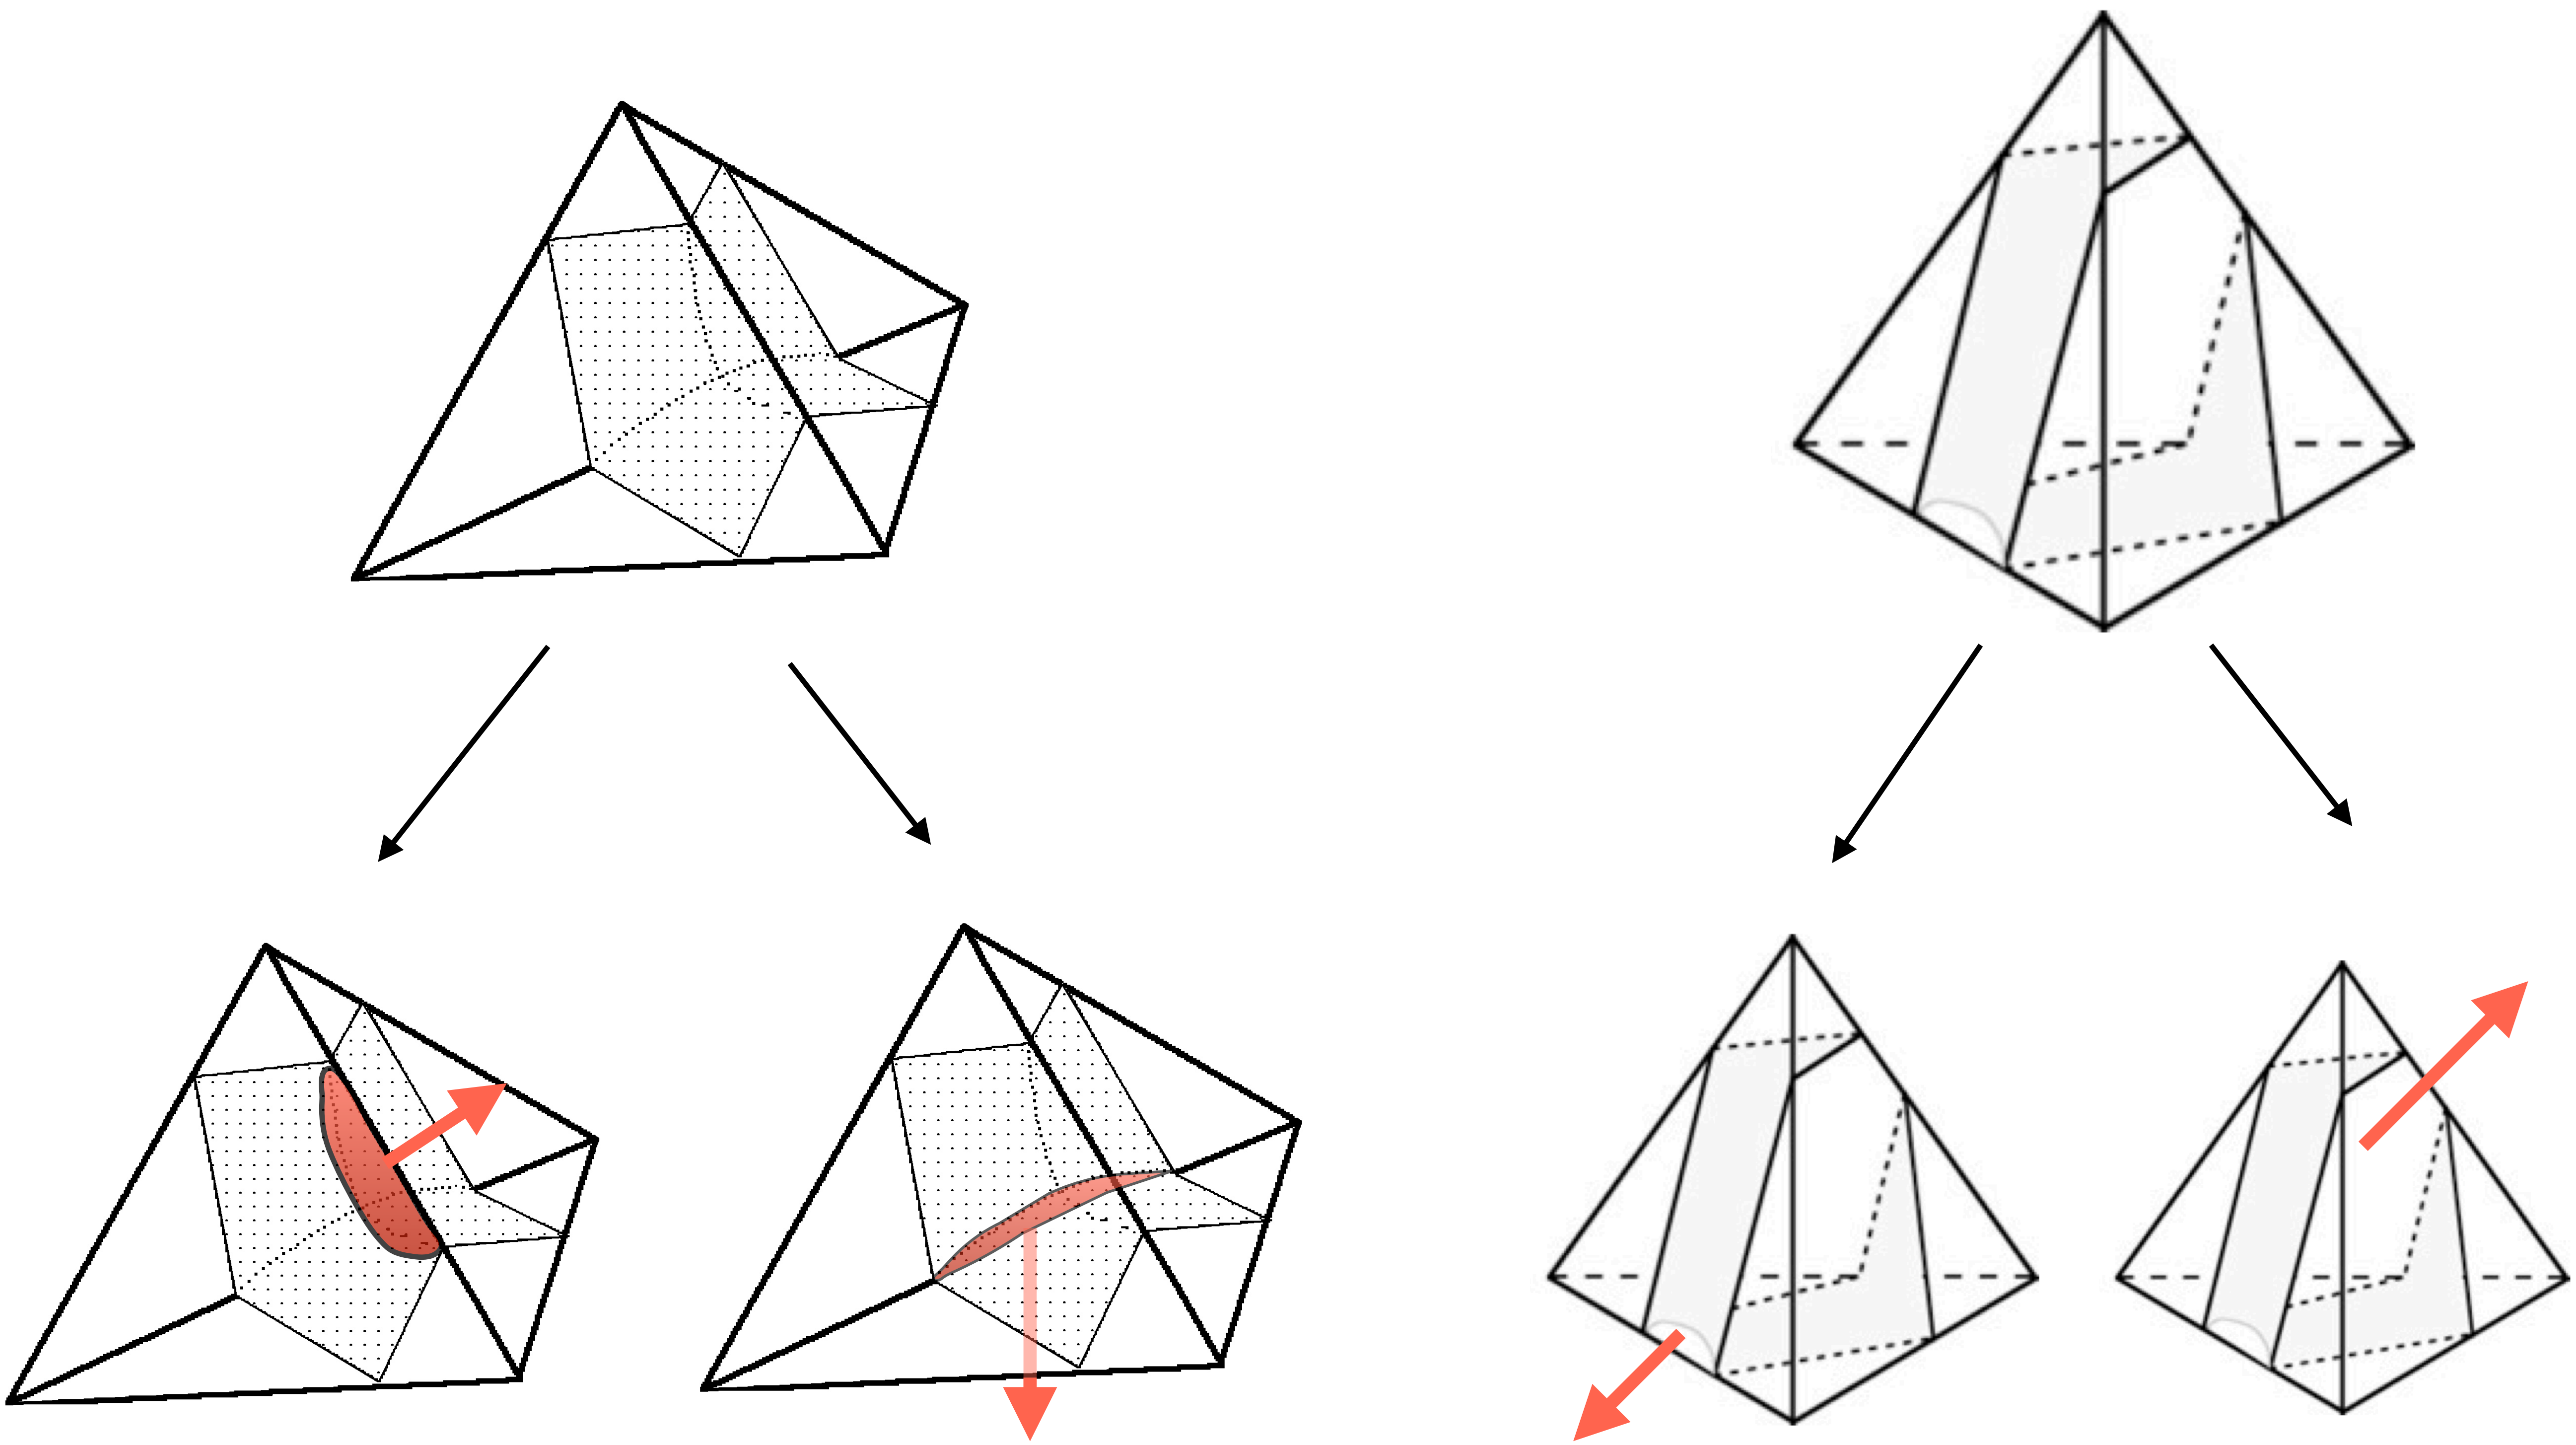

Normalization

Proof. Perform a series of moves on F that

1. Eliminate curves of intersection with a tetrahedron face by compression.
2. Eliminate arcs of intersection with a tetrahedron face by a boundary compression.

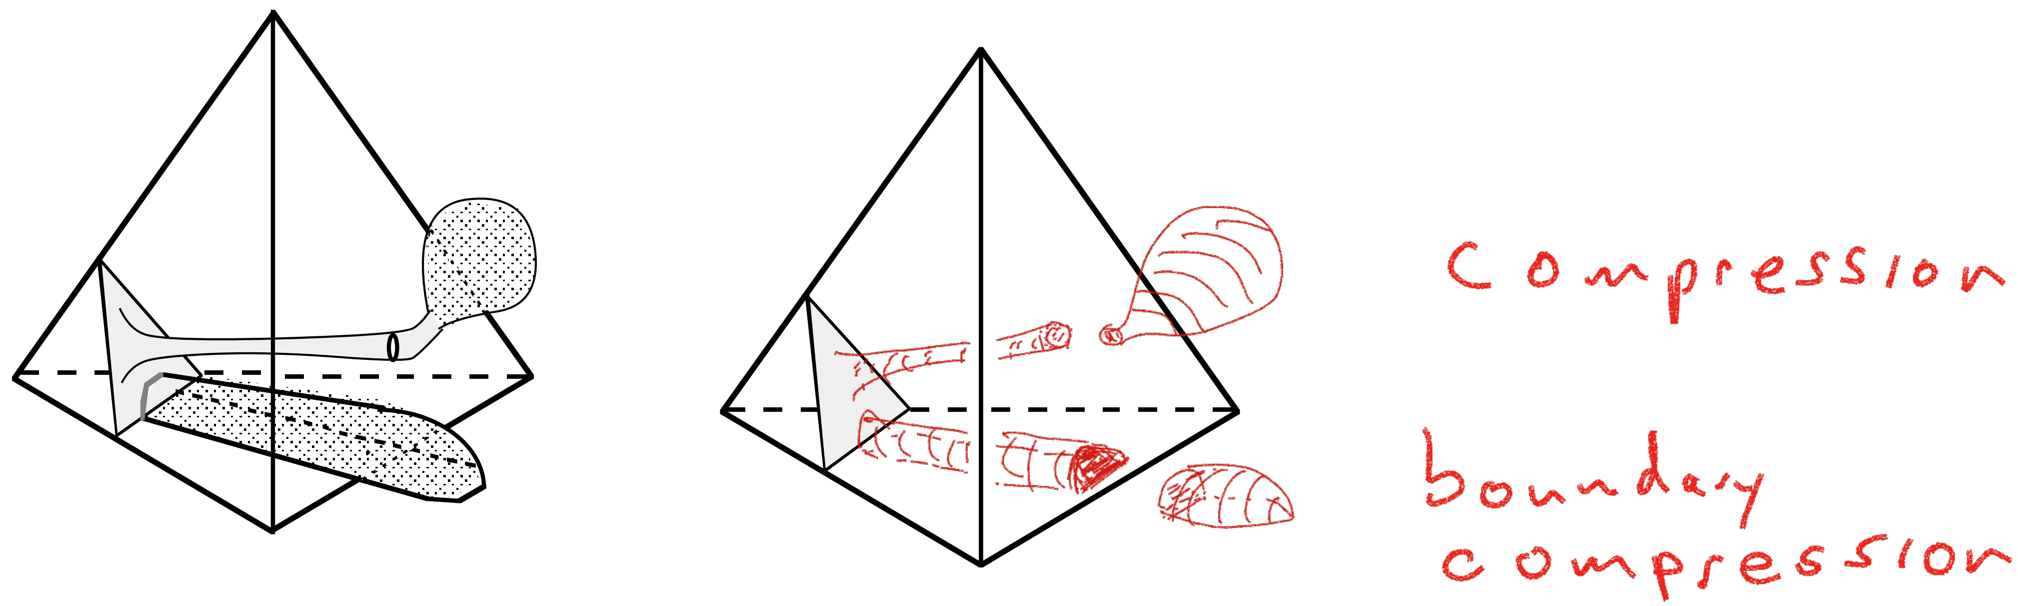

3. Isotop *F* to eliminate pairs of intersection points with a single tetrahedron edge. (If there are several, innermost pair first).

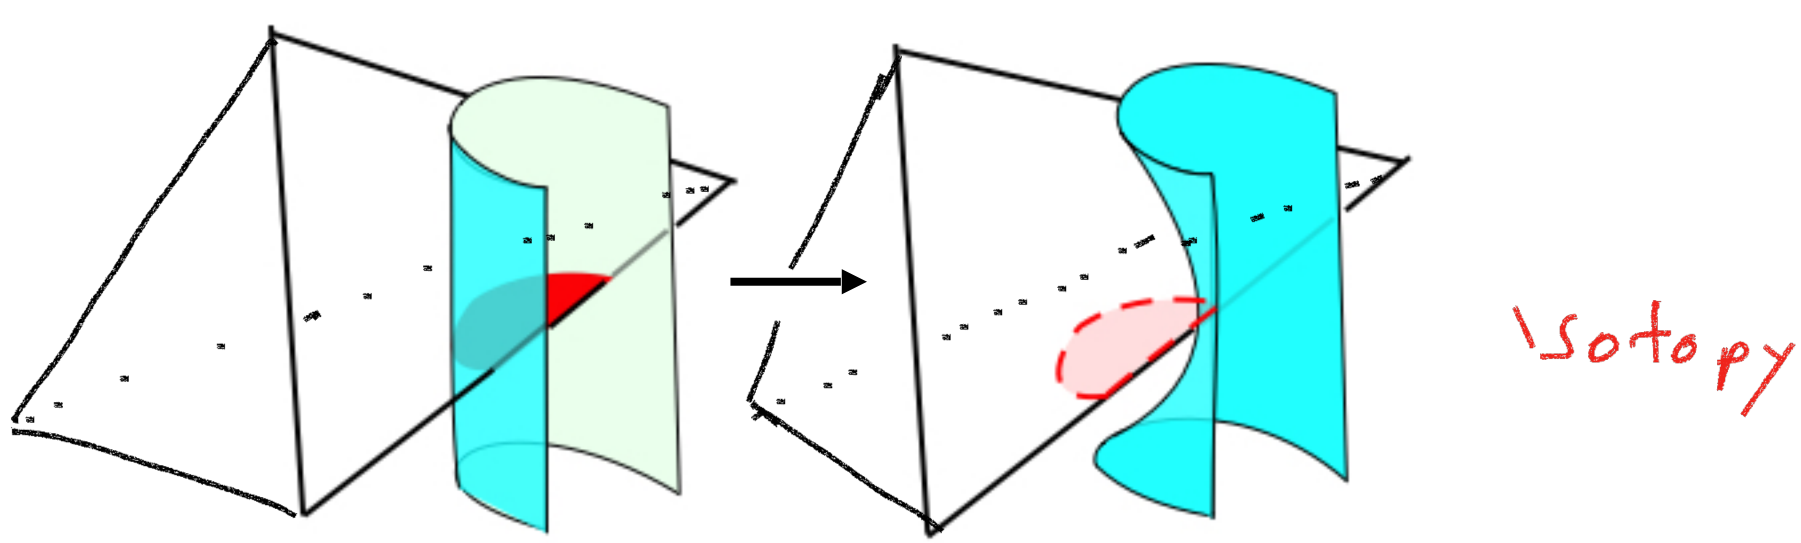

More Normalization

Isotop F to eliminate pairs of intersection points with a single tetrahedron edge. (If there are several, eliminate an innermost pair first)

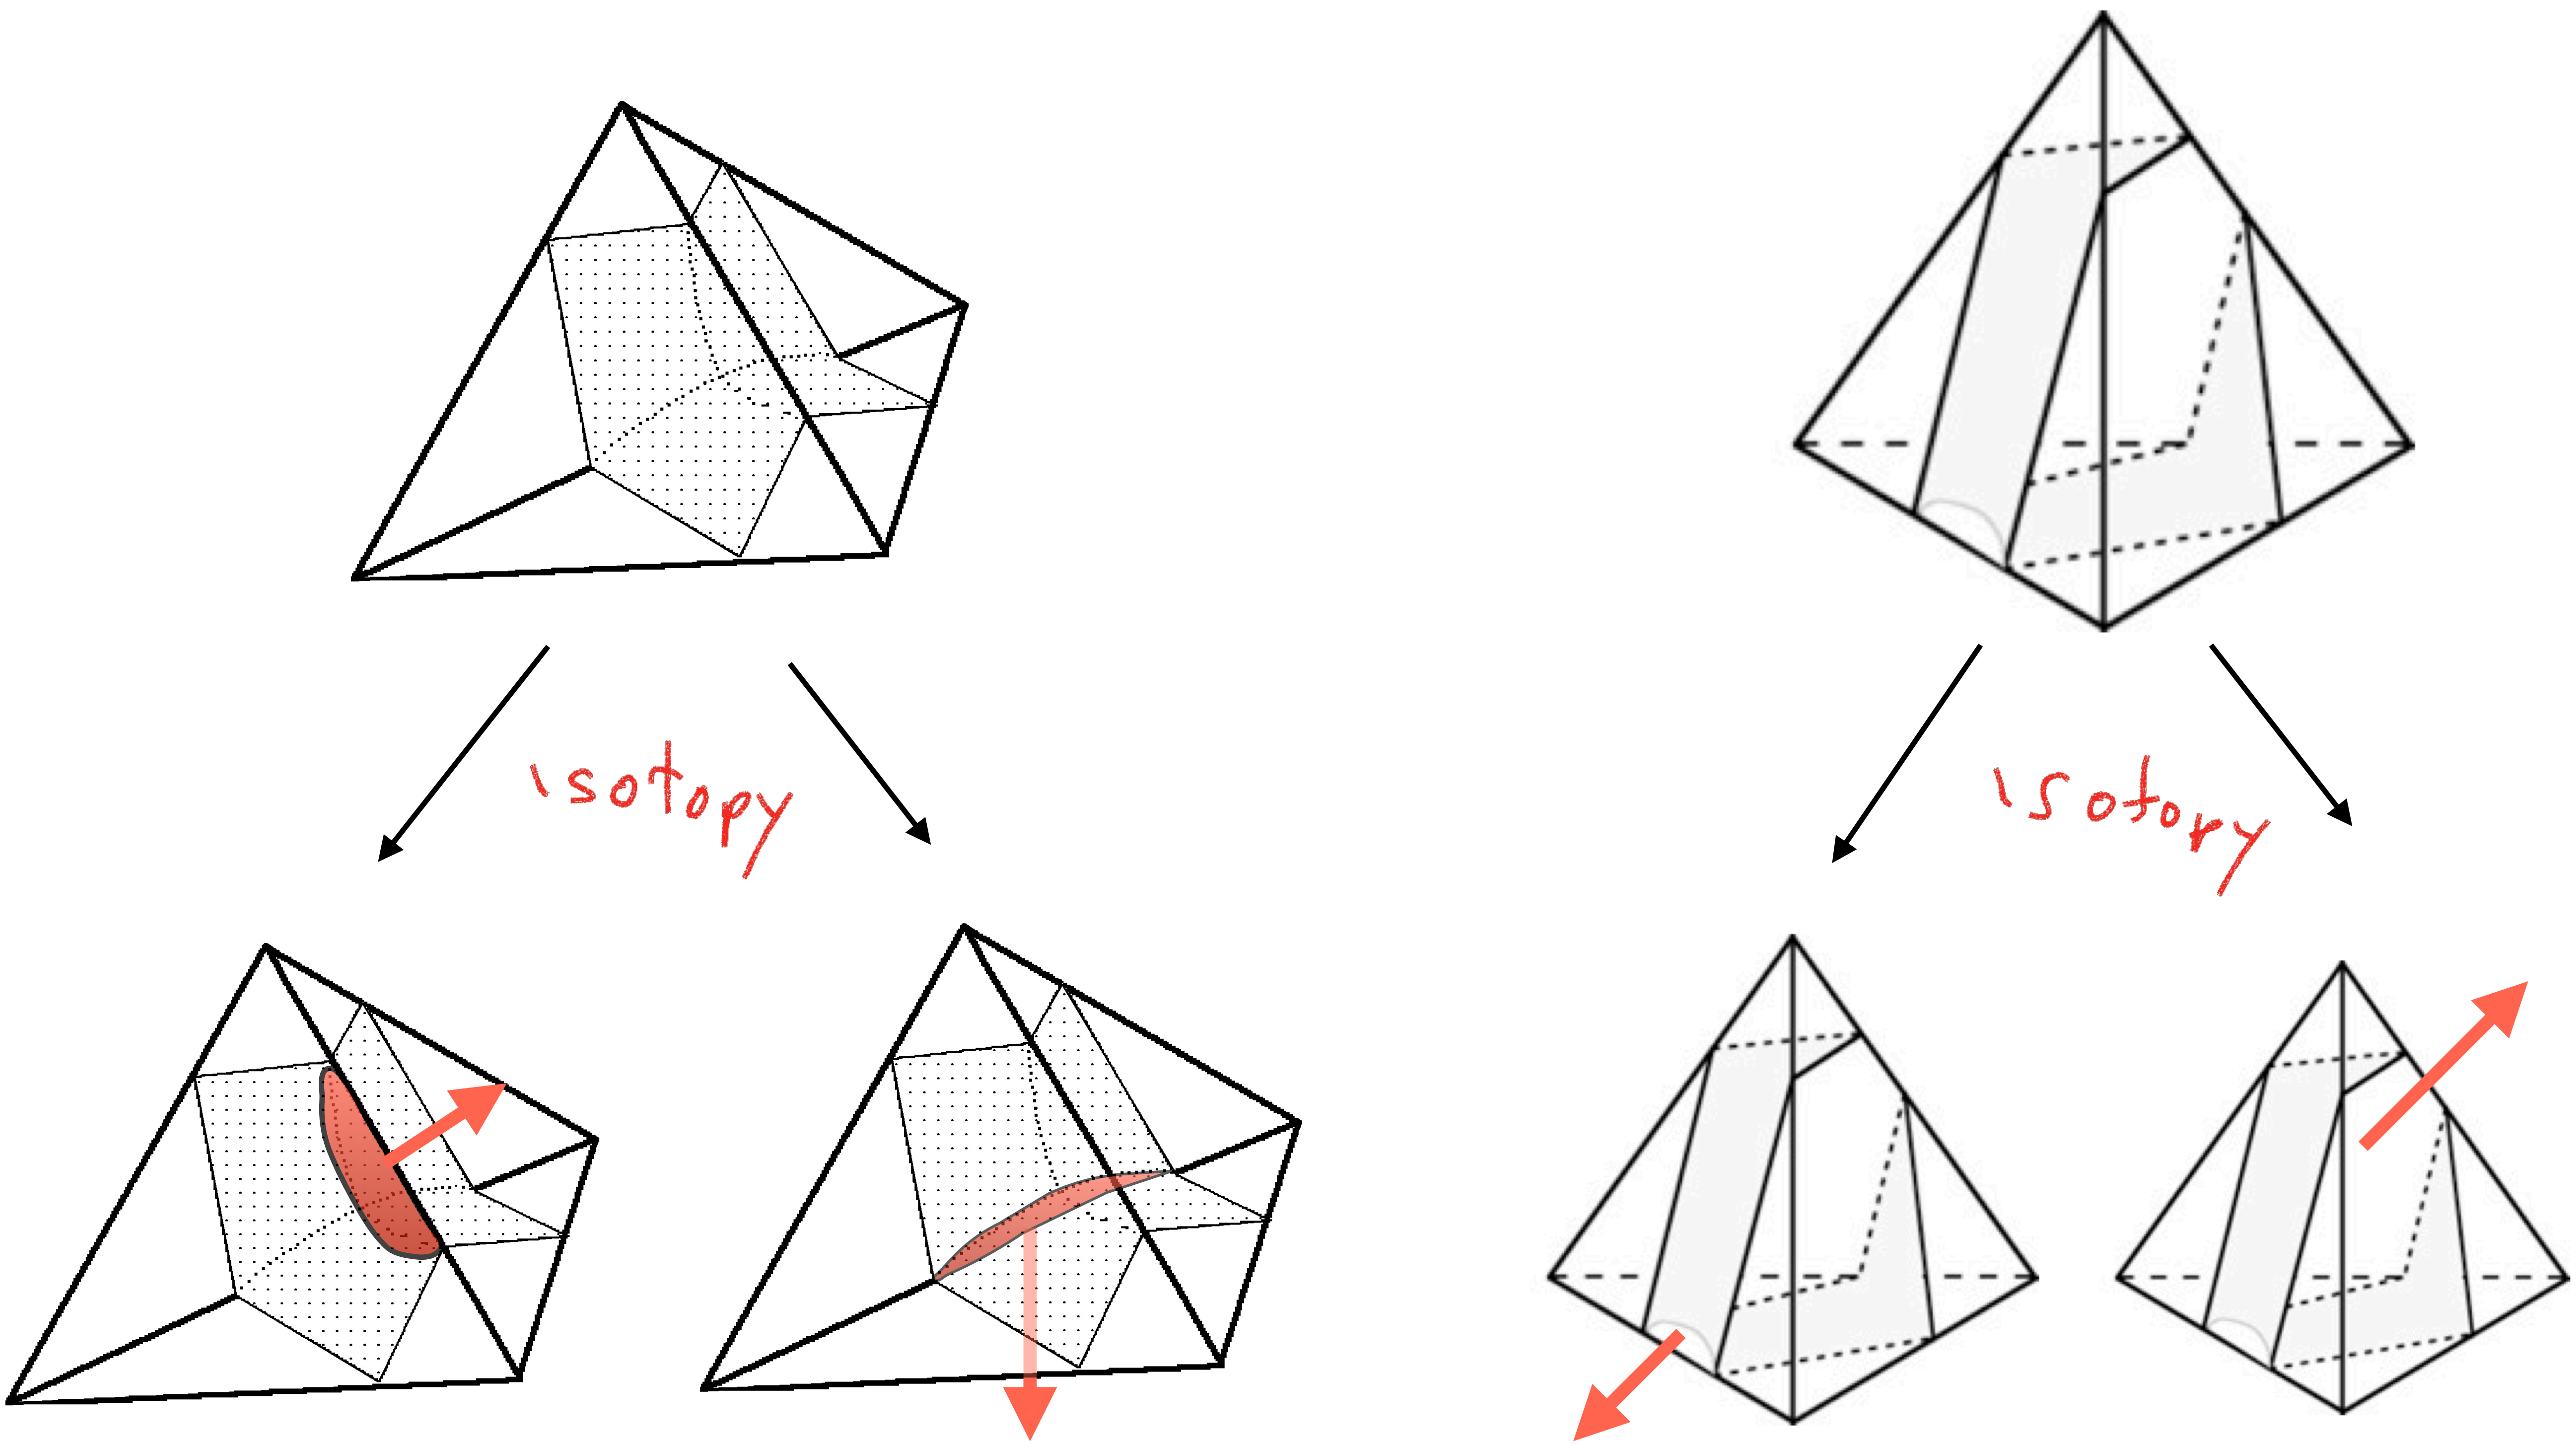

Normalization

Theorem Any surface *F* in a triangulated 3-manifold *M* can be transformed to a normal surface by a sequence of the following moves:
1. Isotopy
2. Compression and boundary-compression
3. Eliminating components lying inside a single tetrahedron.

Proof. Perform a series of simplifying moves on *F*. A connected surface within a tetrahedron that intersects any edge of the tetrahedron in more than one point admits a compression, or boundary compression, or an isotopy that reduces its weight. These moves each reduce the pair (weight, number of curves of intersection of F with faces of the triangulation). So we eventually stop. These moves each reduce the pair (weight, number of curves of intersection of F with faces of the triangulation). So we eventually stop. A connected surface in a tetrahedron that is not compressible and not boundary compressible and which intersects any edge of the tetrahedron in at most one point is a triangle or a quadrilateral.

We end up with a normal surface, (possibly empty).

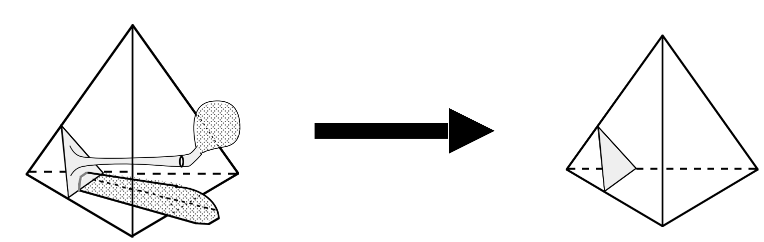

Origins of Normal Surfaces

The problem of understanding essential 2-spheres led to the original definition of normal surfaces by Kneser (1930).

A manifold *M* is *prime* if it is not a connect sum.
This is equivalent to having every 2-sphere in *M* bound a ball,
(except for two cases: $S^1 \times S^2$ and $S^1 \tilde{\times} S^2$.)
Note that S^3 is prime (Alexander's theorem).

Theorem (Kneser)
A closed 3-manifold M has a connect sum decomposition into prime components,

$$M = M_1 \# M_2 \# \ldots \# M_k \qquad \text{with } M_i \text{ prime, } 1 \leq i \leq k.$$

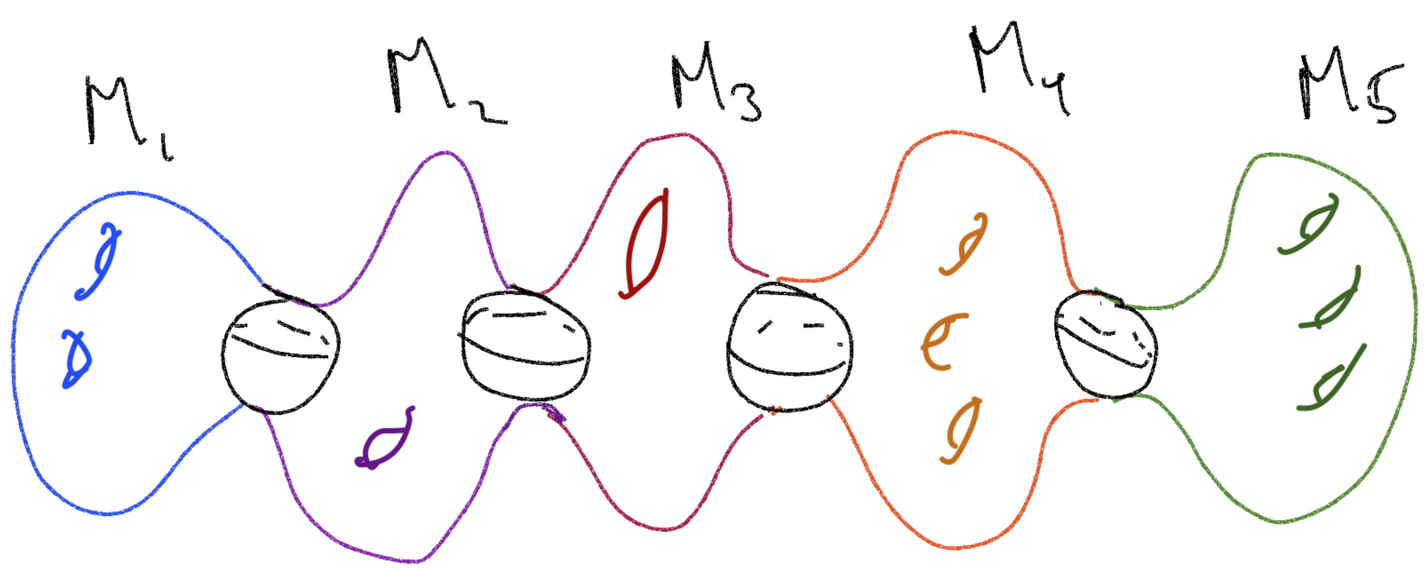

S

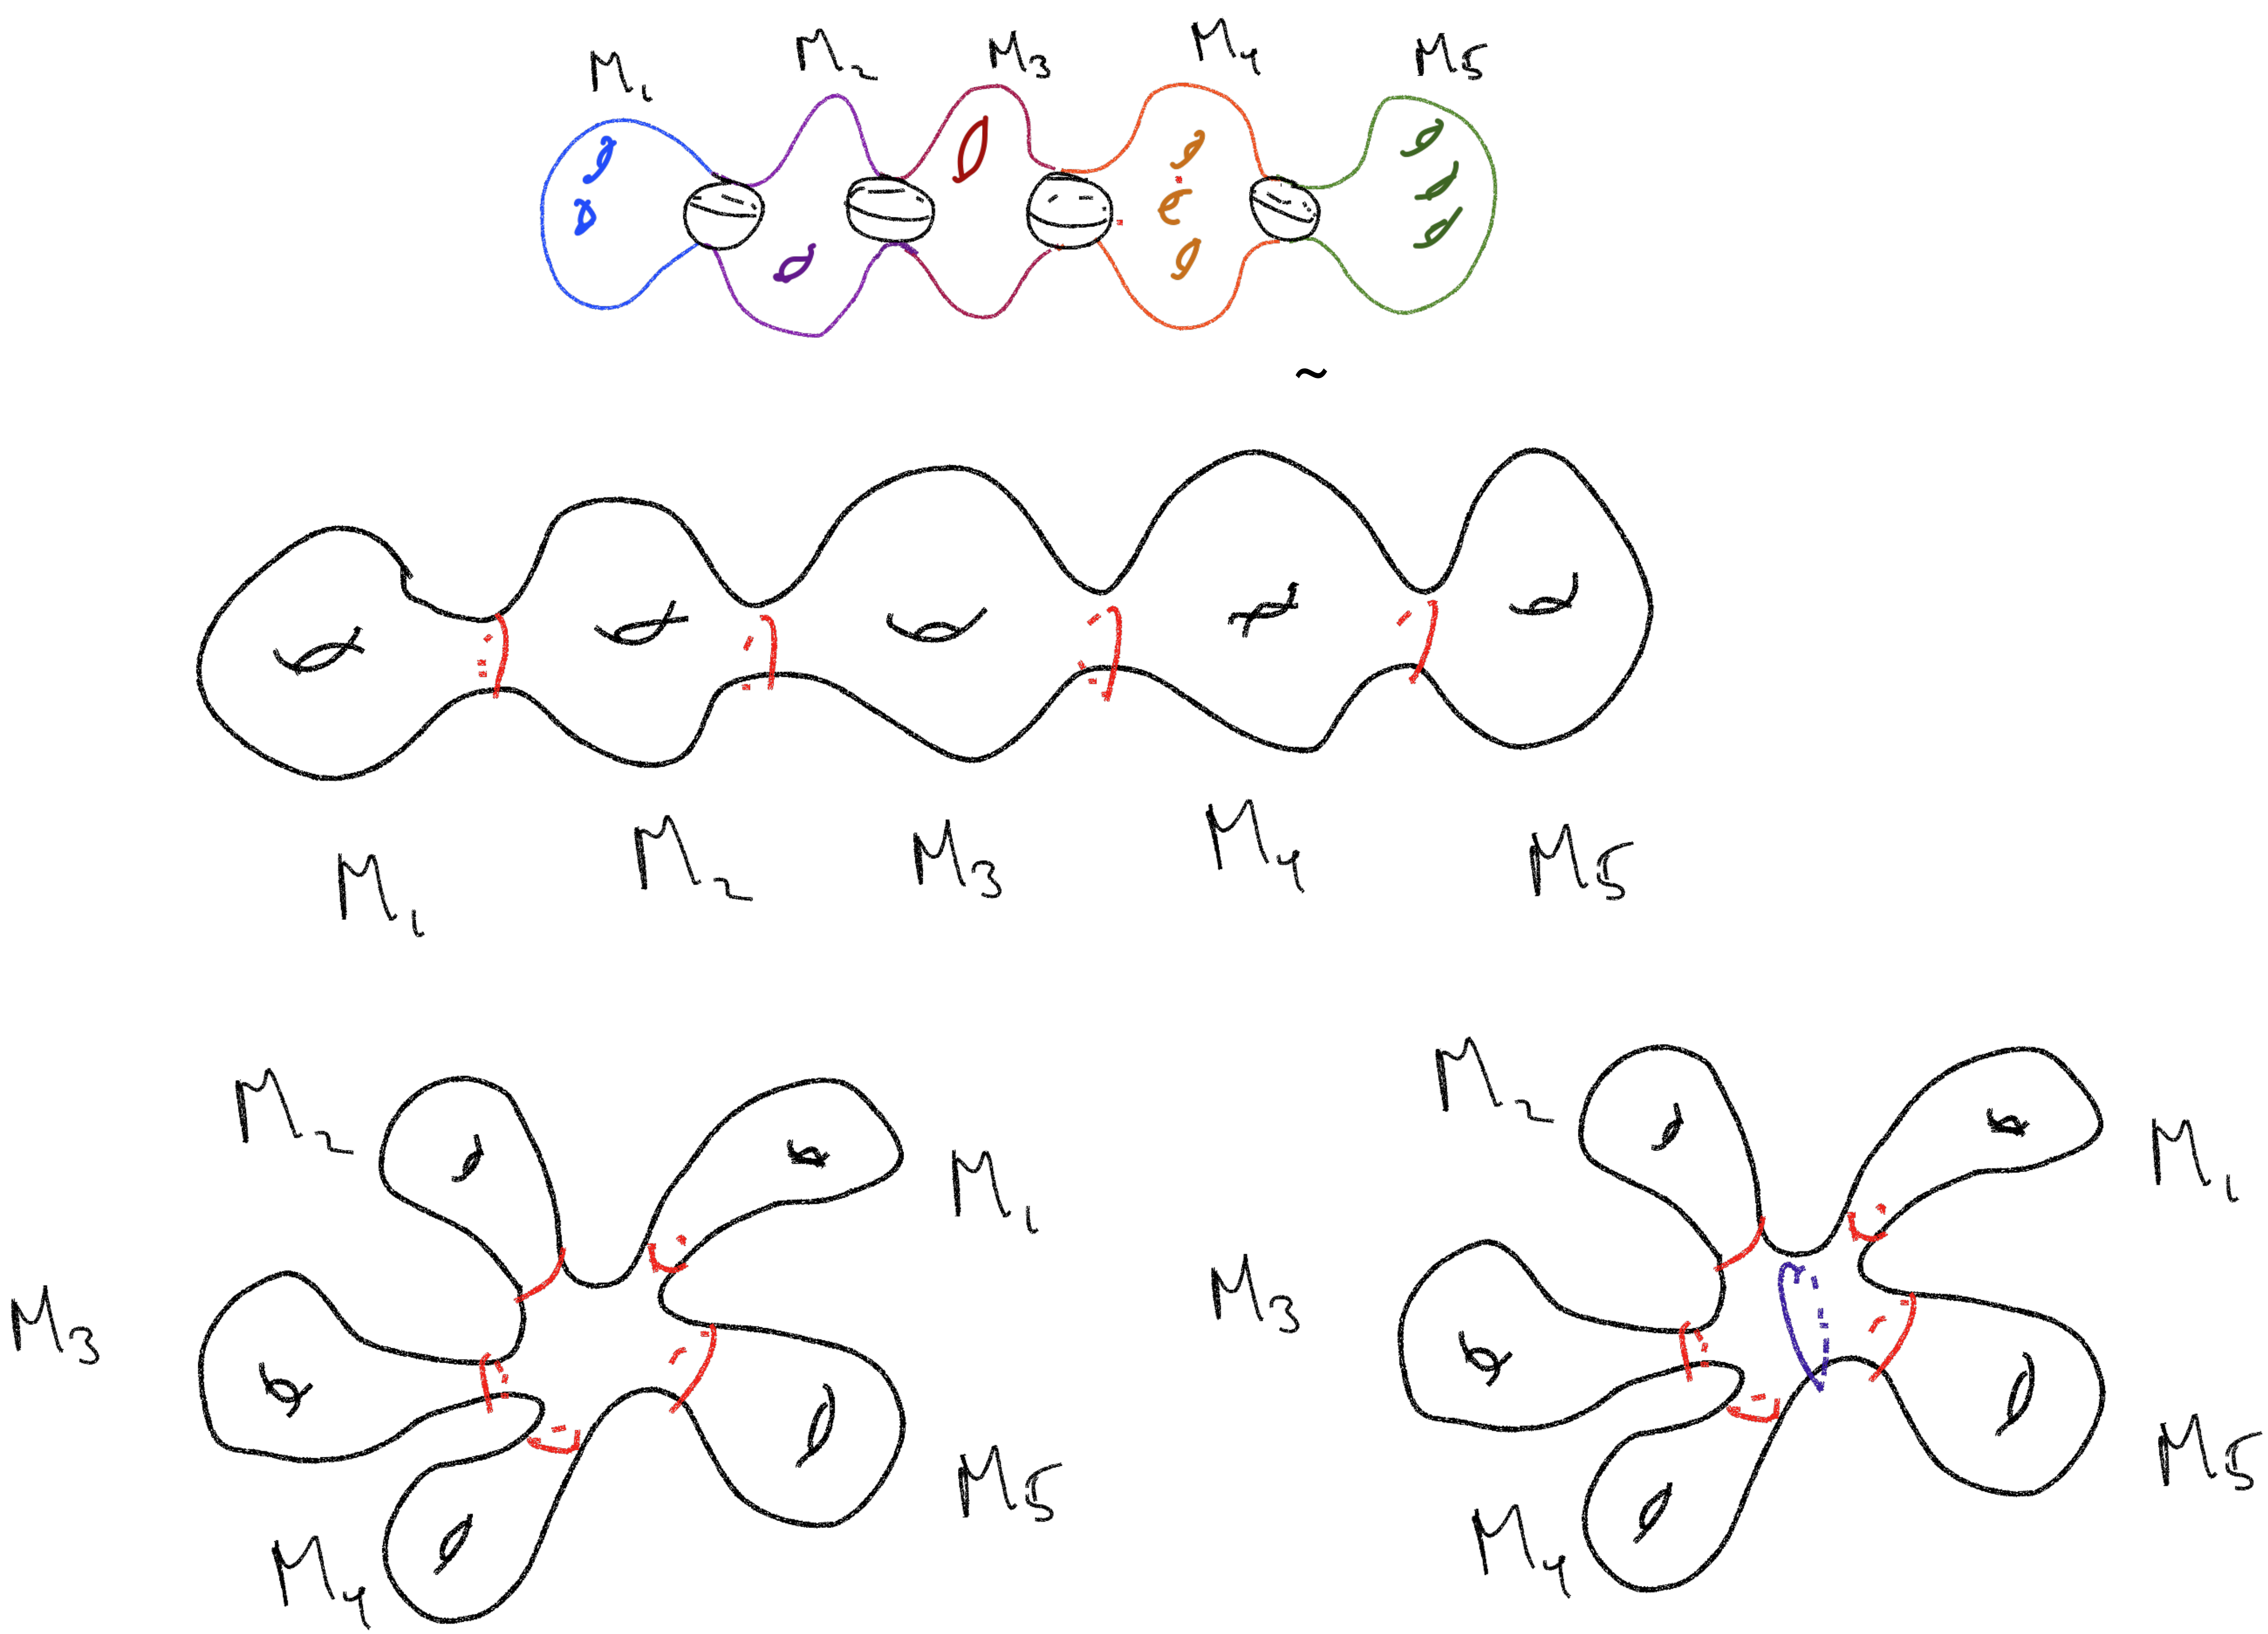

Origins of Normal Surfaces

The problem of understanding essential 2-spheres led to the original definition of normal surfaces by Kneser (1930).

A manifold M is *prime* if it is not a connect sum.
This is equivalent to having every 2-sphere in M bound a ball, (except for two cases: $S^1 \times S^2$ and $S^1 \tilde{\times} S^2$.) Note that S^3 is prime (Alexander's theorem).

Theorem (Kneser)
A closed 3-manifold M has a connect sum decomposition into prime components,

$$M = M_1 \# M_2 \# \ldots \# M_k \qquad \text{with } M_i \text{ prime, } 1 \leq i \leq k.$$

Origins of Normal Surfaces

The problem of understanding essential 2-spheres led to the original definition of normal surfaces by Kneser (1930).

A manifold *M* is *prime* if it is not a connect sum.
This is equivalent to having every 2-sphere in *M* bound a ball, (except for two cases: $S^1 \times S^2$ and $S^1 \tilde{\times} S^2$.) Note that S^3 is prime (Alexander's theorem).

Theorem (Kneser)
A closed 3-manifold M has a connect sum decomposition into prime components,

$$M = M_1 \# M_2 \# \ldots \# M_k \qquad \text{with } M_i \text{ prime, } 1 \leq i \leq k.$$

Proof A set of disjoint 2-spheres is *independent* if no pair are parallel, and no subset bounds a punctured ball. A connect sum decomposition gives *k*-1 independent 2-spheres.

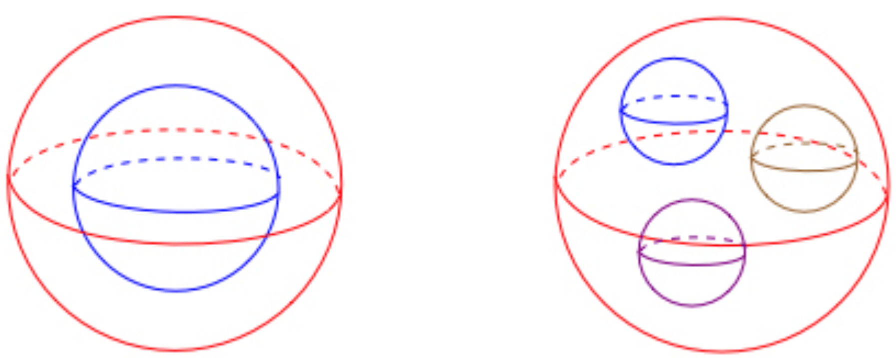

The 2-spheres in a connect sum decomposition are independent.

Kneser's Theorem follows when we answer:

Question: How many 2-spheres can there be in an independent set?

Lemma 1: In an oriented *M* with *t* 3-simplices, the size of an independent set of 2-spheres is at most *7t*.

So the number *k* of prime summands is at most 7t-1.

Counting Prime Components

Lemma 1: In a 3-manifold *M* that is triangulated with *t* 3-simplices, the size of an independent set of 2-spheres is at most *7t*.

Follows from:

Lemma 2: In an oriented 3-manifold *M* triangulated with *t* 3-simplices, the number of disjoint, orientable normal 2-spheres is at most *7t*.

and

Lemma 3: If there is an independent set of *k* 2-spheres, then there is an independent set of *k* normal 2-spheres.

Counting Disjoint Normal Surfaces

Lemma 2: In an oriented 3-manifold *M* triangulated with *t* 3-simplices, the number of disjoint, orientable normal 2-spheres is at most *7t*.

Proof: Let *G* be a maximal family of disjoint normal 2-spheres, with no pair parallel. Some of these (at most *t*) are vertex linking 2-spheres.

Each tetrahedron *T* meets *G* in a collection of triangles and quadrilaterals. The complement *T*-*G* consists of a collection of product regions (triangle x I and quadrilateral x I) and at most two other "bad" regions not meeting a vertex. The total number of bad regions is at most 2*t*.

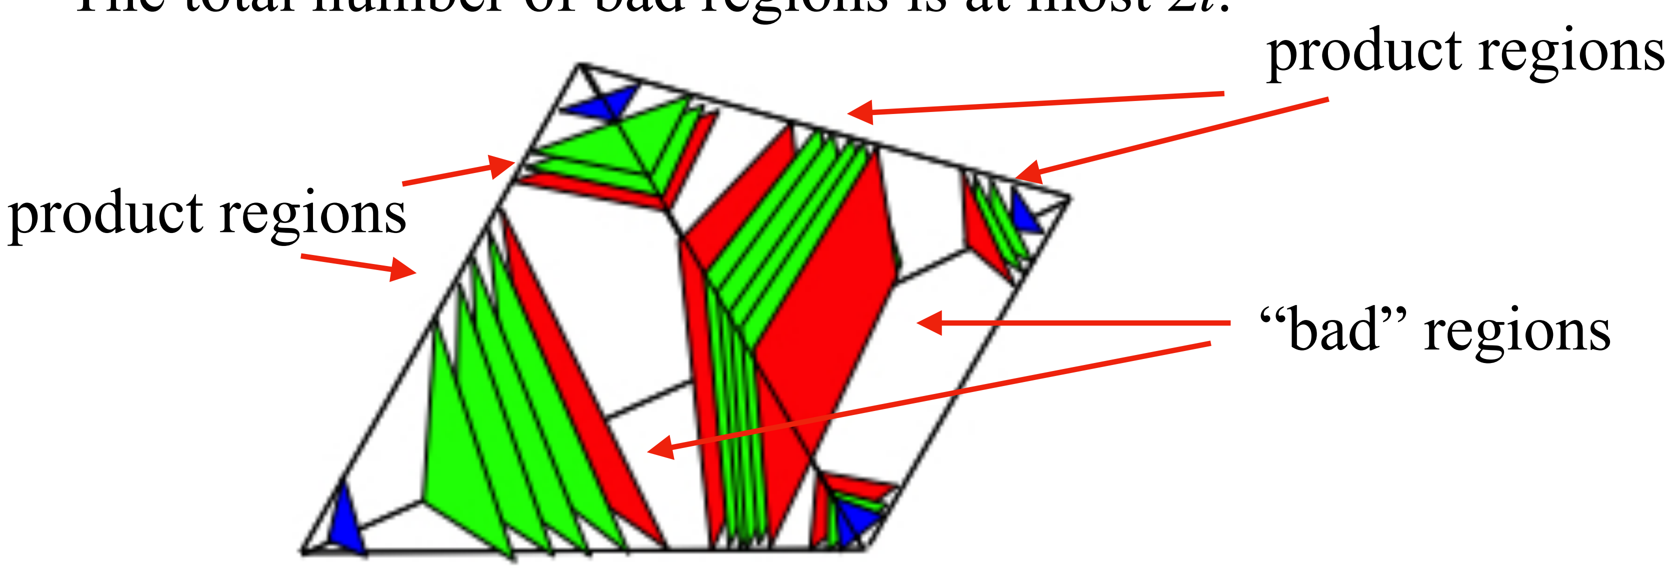

6t red faces in M.

Each oriented normal sphere in *G* meets at least one of the bad regions (contains a red face), except perhaps for vertex linking 2-spheres. Otherwise the sphere would be parallel to another sphere on one side or the other.

The bad regions have a total of at most 6*t* red faces. So there are at most 6t spheres in addition to at most *t* vertex-linking 2-spheres, for a total of 7t.

Counting Disjoint Normal Surfaces

Lemma 3: If there is an independent set of n 2-spheres in M, then there is an independent set of n normal 2-spheres.

Proof: The collection of 2-spheres can be normalized by isotopy and compression.

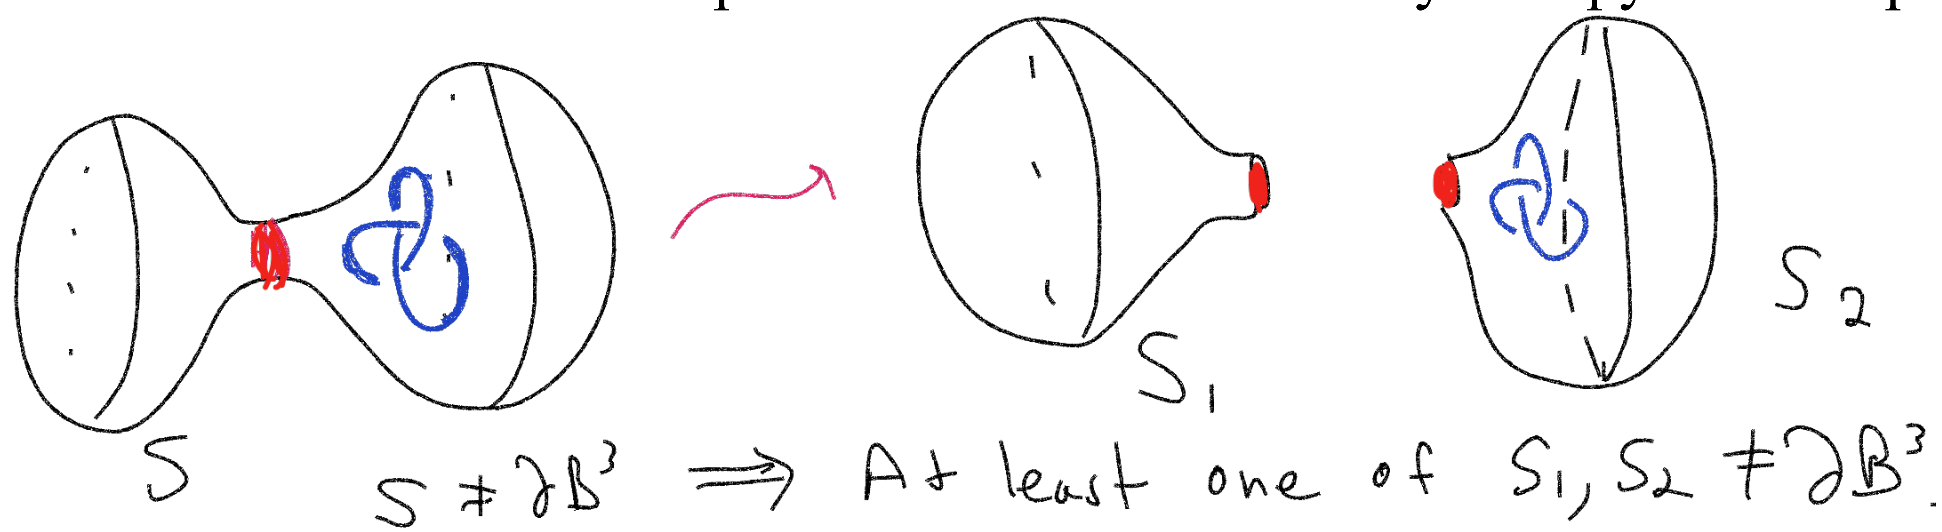

Normalizing does not decrease the number of independent 2-spheres.

Isotopy: Just moves things around.

Compression: Compressing a 2-sphere can change the independent set, but does not decrease the number of 2-spheres in it, since the union of two punctured 3-balls along a disk is a punctured 3-ball. If both S_1 and S_2 are dependent, then so is S.

Theorem (Kneser)
A closed 3-manifold M has a connect sum decomposition into prime components:
$M = M_1 \# M_2 \# \ldots \# M_k$ with M_i prime, $1 \leq i \leq k$.

Normal Surfaces in Algorithms

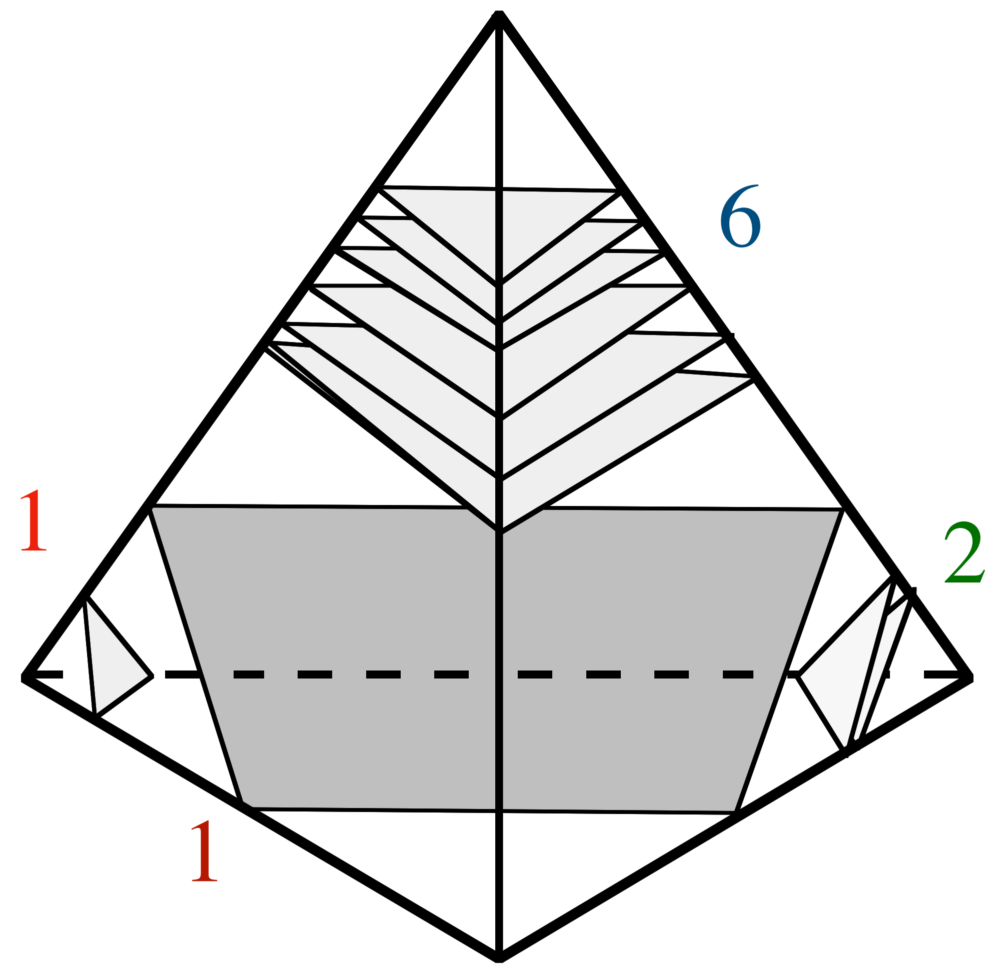

Unknotting

Triangulate a tetrahedron containing a knot *K* so that the knot lies on the edges. This gives a finite combinatorial description.

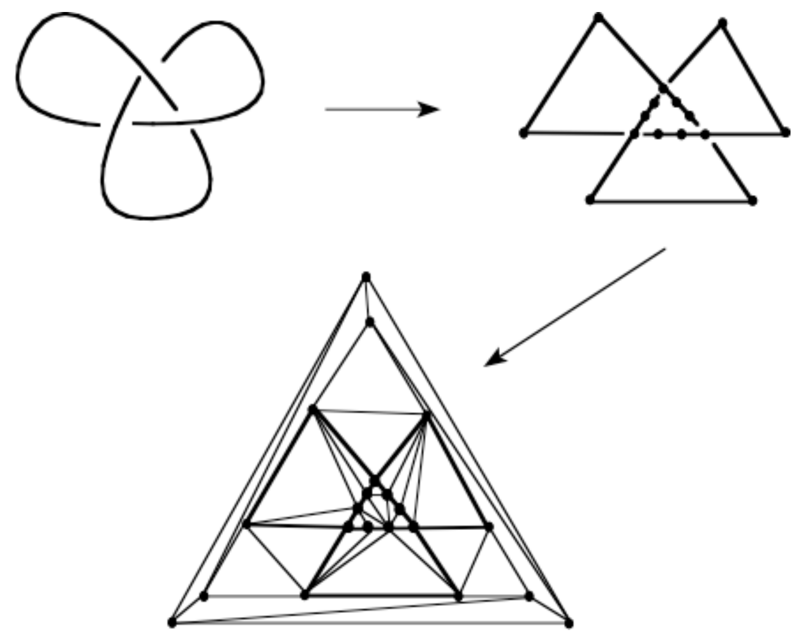

Subdivide this triangulation and remove a regular neighborhood of *K*. This gives a 3-manifold homeomorphic to the knot complement, $S^3 - K$ (or $S^3 \backslash K$). This manifold M_K has a torus boundary.

$M_K = \overline{S^3 \backslash K}$ **=**  ***??***

An Unknotting Algorithm

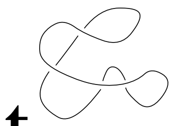

Triangulate a tetrahedron containing a knot K so that the knot lies on the edges. This gives a finite combinatorial description.

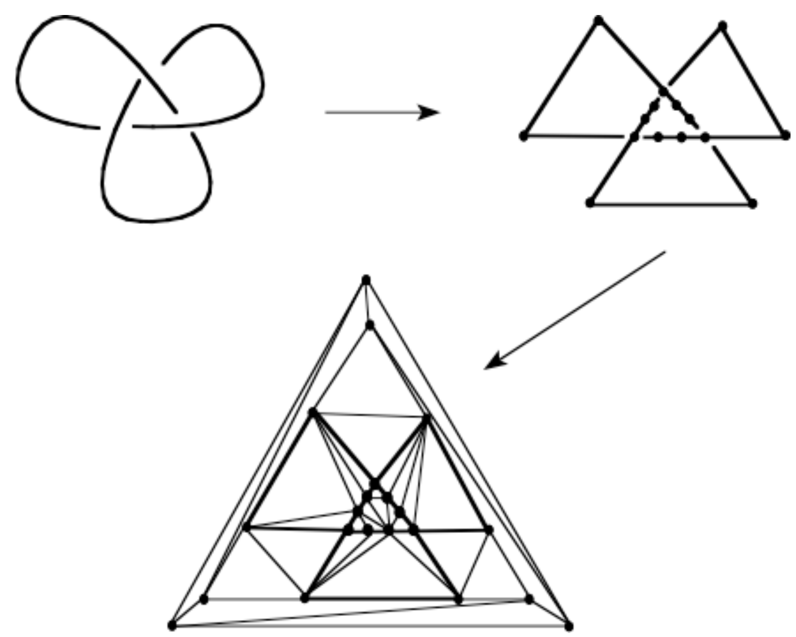

Subdivide this triangulation and remove a regular neighborhood of K. This gives a 3-manifold homeomorphic to the knot complement, M_K (or $\overline{S^3 \setminus K}$). This manifold M_K has a torus boundary.

$$\boldsymbol{M_K = \overline{S^3 \setminus K}} \quad = $$

??

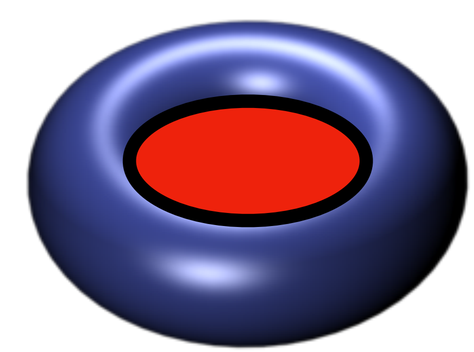

The knot K is unknotted if and only if this torus is compressible.

Normal Surfaces and Knots

Question: Starting from a knot with n-crossings, how many tetrahedra are needed to form a simplicial complex with the knot embedded on its edges?

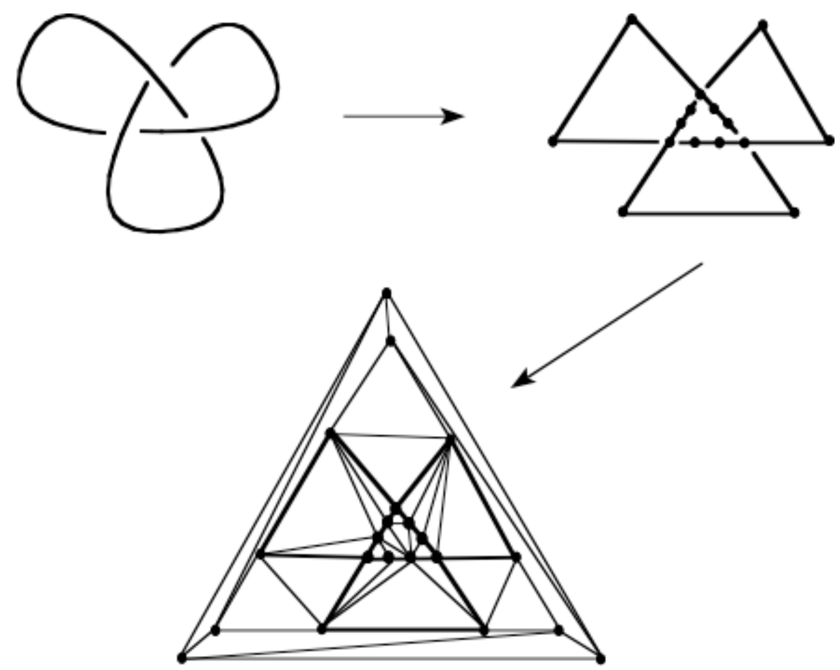

This type of question becomes important when we study the running time (computational complexity) of algorithms.
For now we just need to observe the number can be bounded explicitly with some work.

Normal Surfaces and Knots

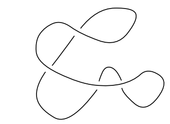

Question: Starting from a knot with n-crossings, how many tetrahedra are needed to form a simplicial complex with the knot embedded on its edges?

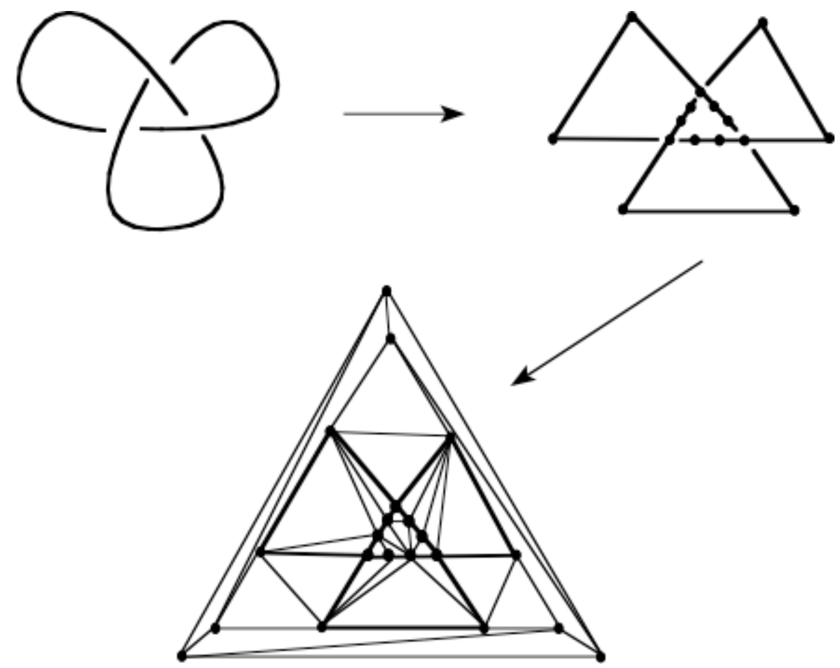

Lemma (H-Lagarias-Pippenger, 1999)
Given a knot diagram D with n crossings, one can construct in time $O(n \log n)$ a combinatorial triangulation of S^3 using at most 253,440(n+ 1) tetrahedra, which contains a good triangulation of M_K.

Normal Surface to Vector

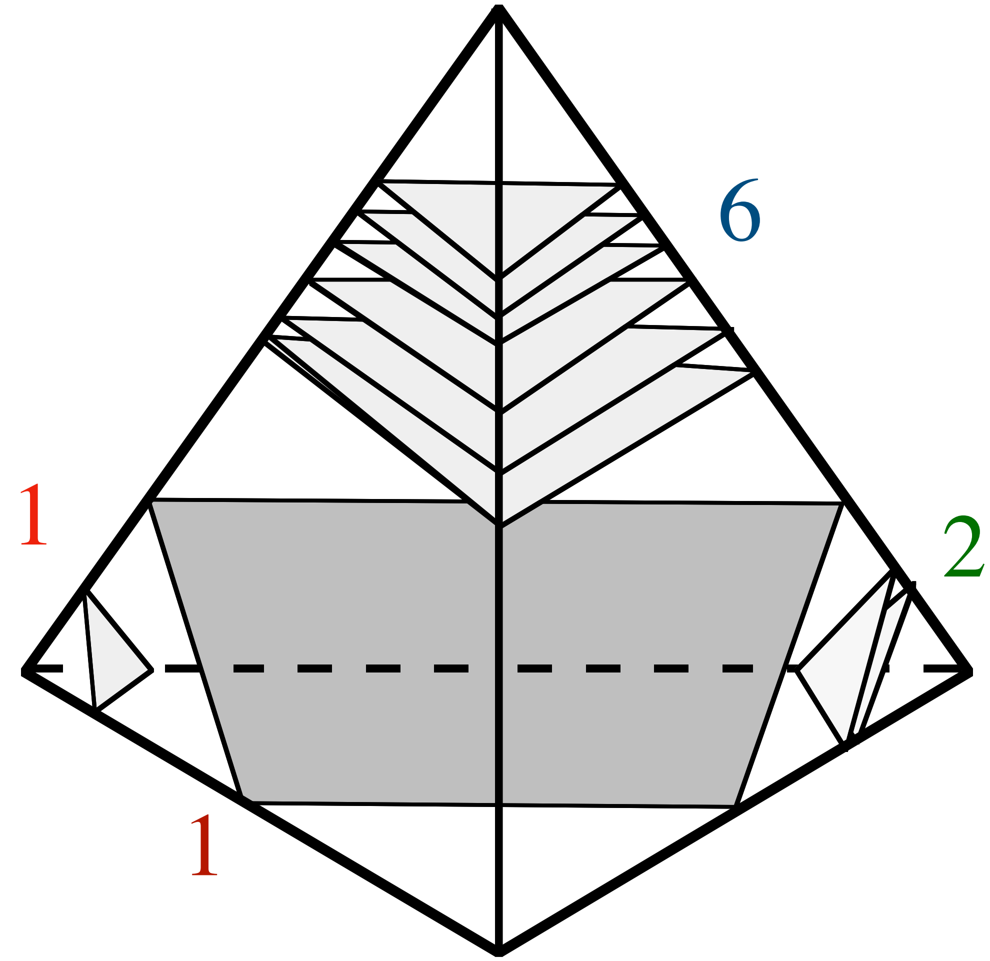

In each tetrahedron, a normal surface is encoded by 7 integers. These count the number of each of the four triangle types and 3 quadrilateral types. The above surface produces a vector with 7 entries:

$$(1, 6, 2, 0, 1, 0, 0)$$

Repeating for each of *t* tetrahedra gives 7t integers.

Efficient Representations

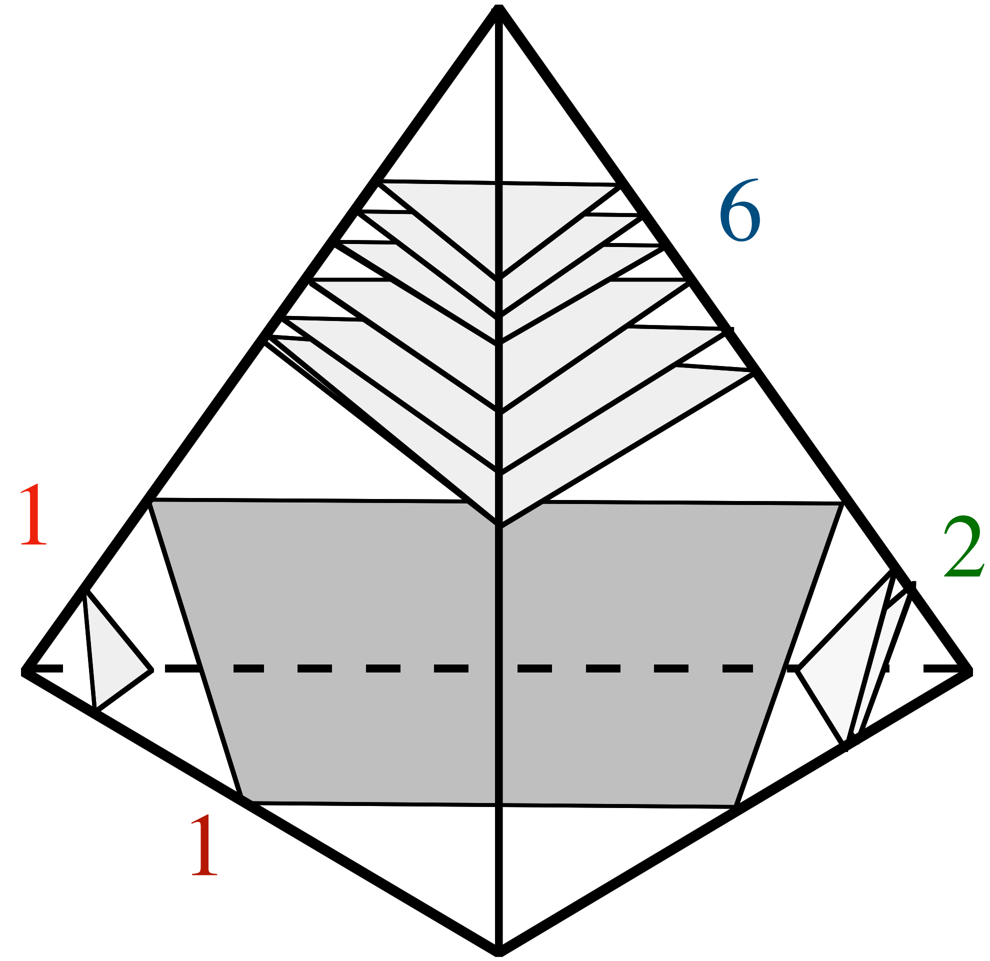

Normal surfaces give very efficient descriptions of a surface.
Exponentially better than “standard” descriptions.

$$(1, 6, 2, 0, 1, 0, 0)$$

The number of vertices in the above surface grows exponentially faster than the size of this vector.
This gives a very efficient surface description.

Normal Disks for Knots

A spanning disk for a knot in M_K is described by a vector.

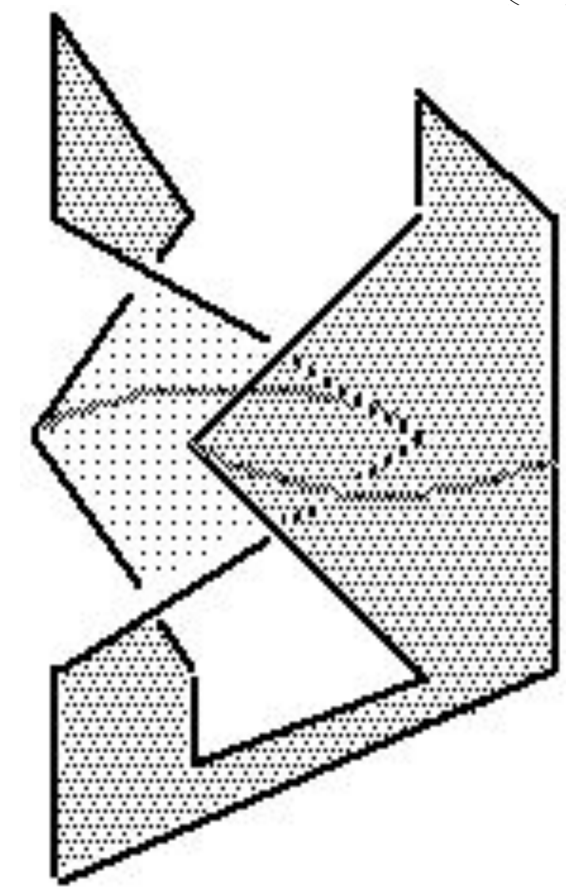

$(v_1, v_2, v_3, \ldots, v_{7t})$

These vectors are not arbitrary.

1. They have integer entries
2. Each entry is non-negative
3. The vectors for two adjoining tetrahedra match up along their common face.

Matching Equations

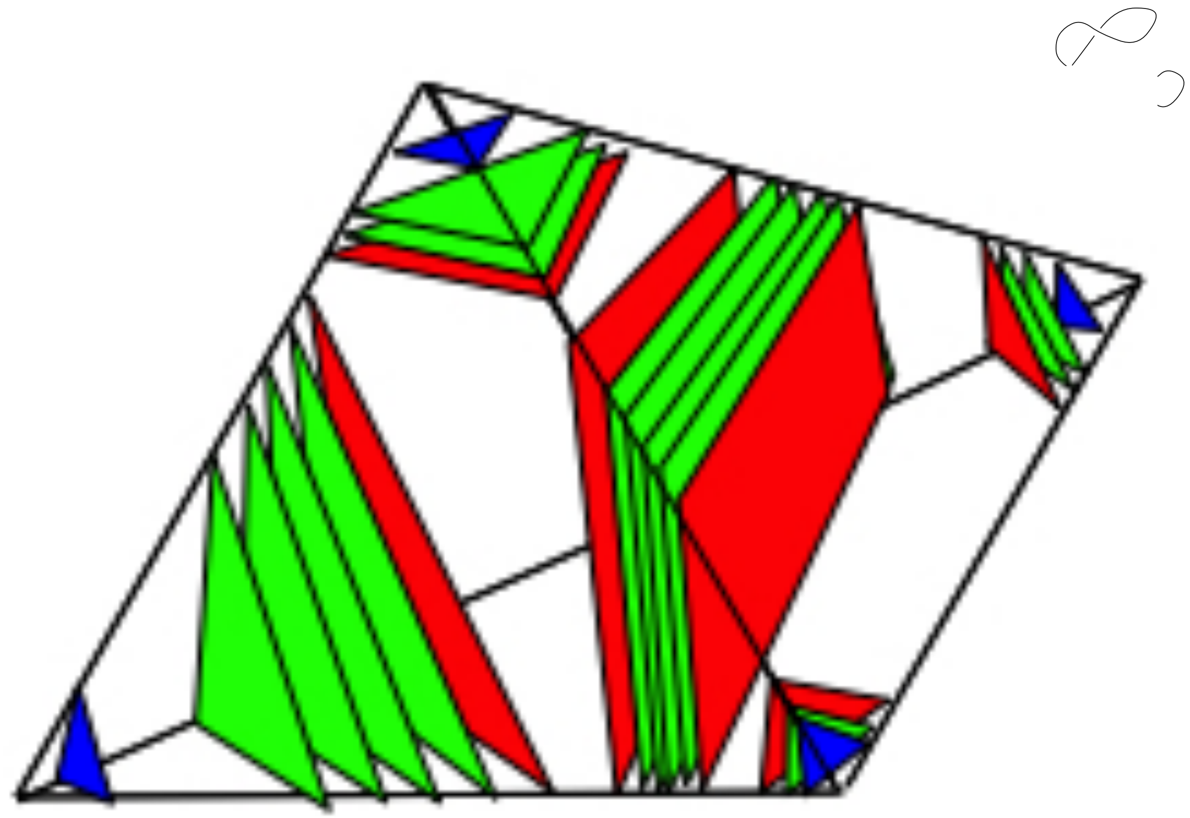

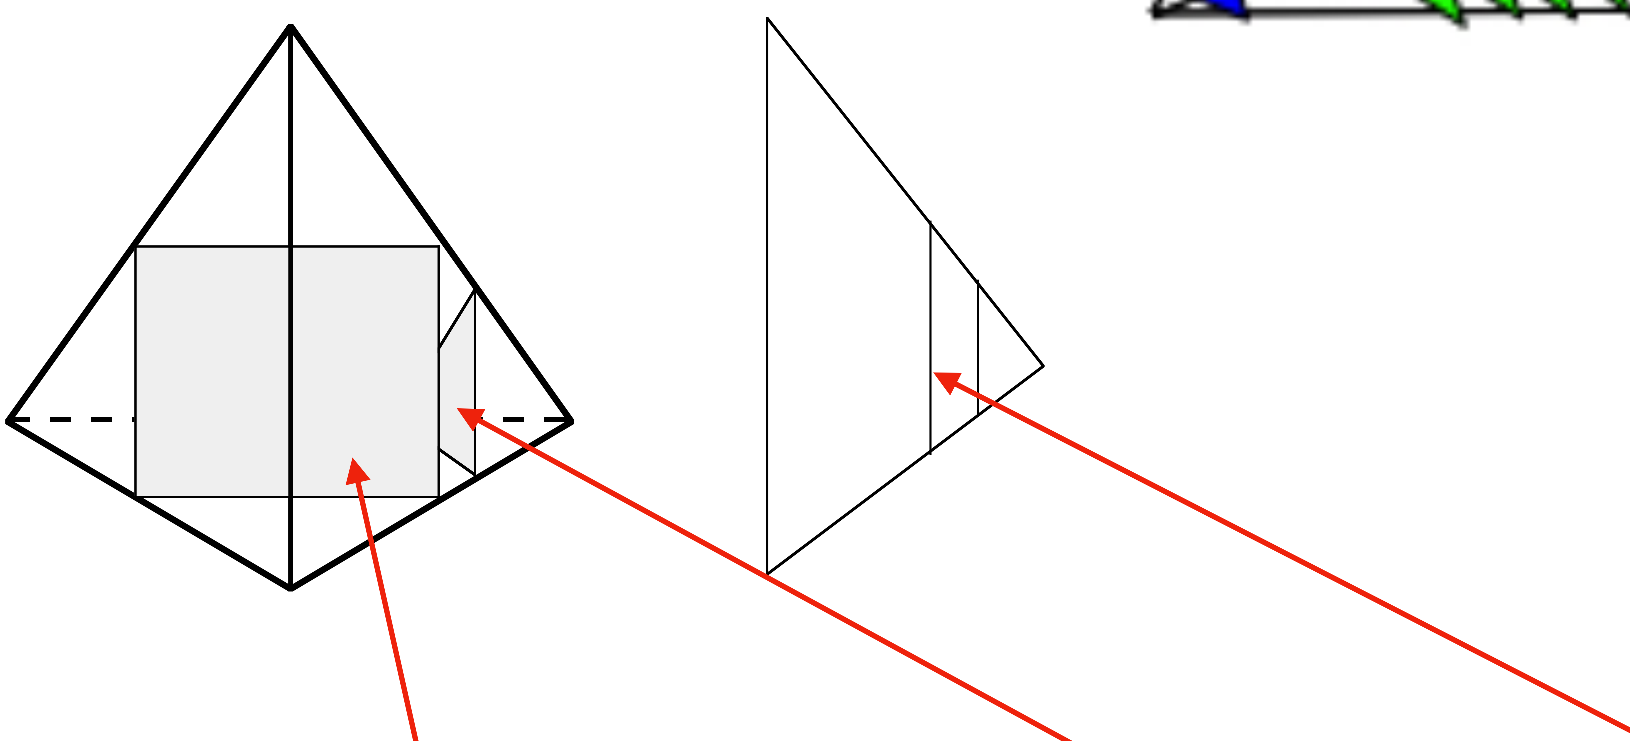

One type of quadrilateral and one type of triangle contribute to edges of this type on this face of a tetrahedron. The number of such edges matches along adjacent tetrahedra.

Matching Equations

To fit together to give a surface, pieces must match up across tetrahedra with common faces.

This gives linear equations for the integer coordinates of the vector $(v_1, v_2, v_3, \ldots\ v_{7t})$ of the form:

$$v_i+v_j = v_k+v_l \qquad \text{with} \qquad v_i \geq 0.$$

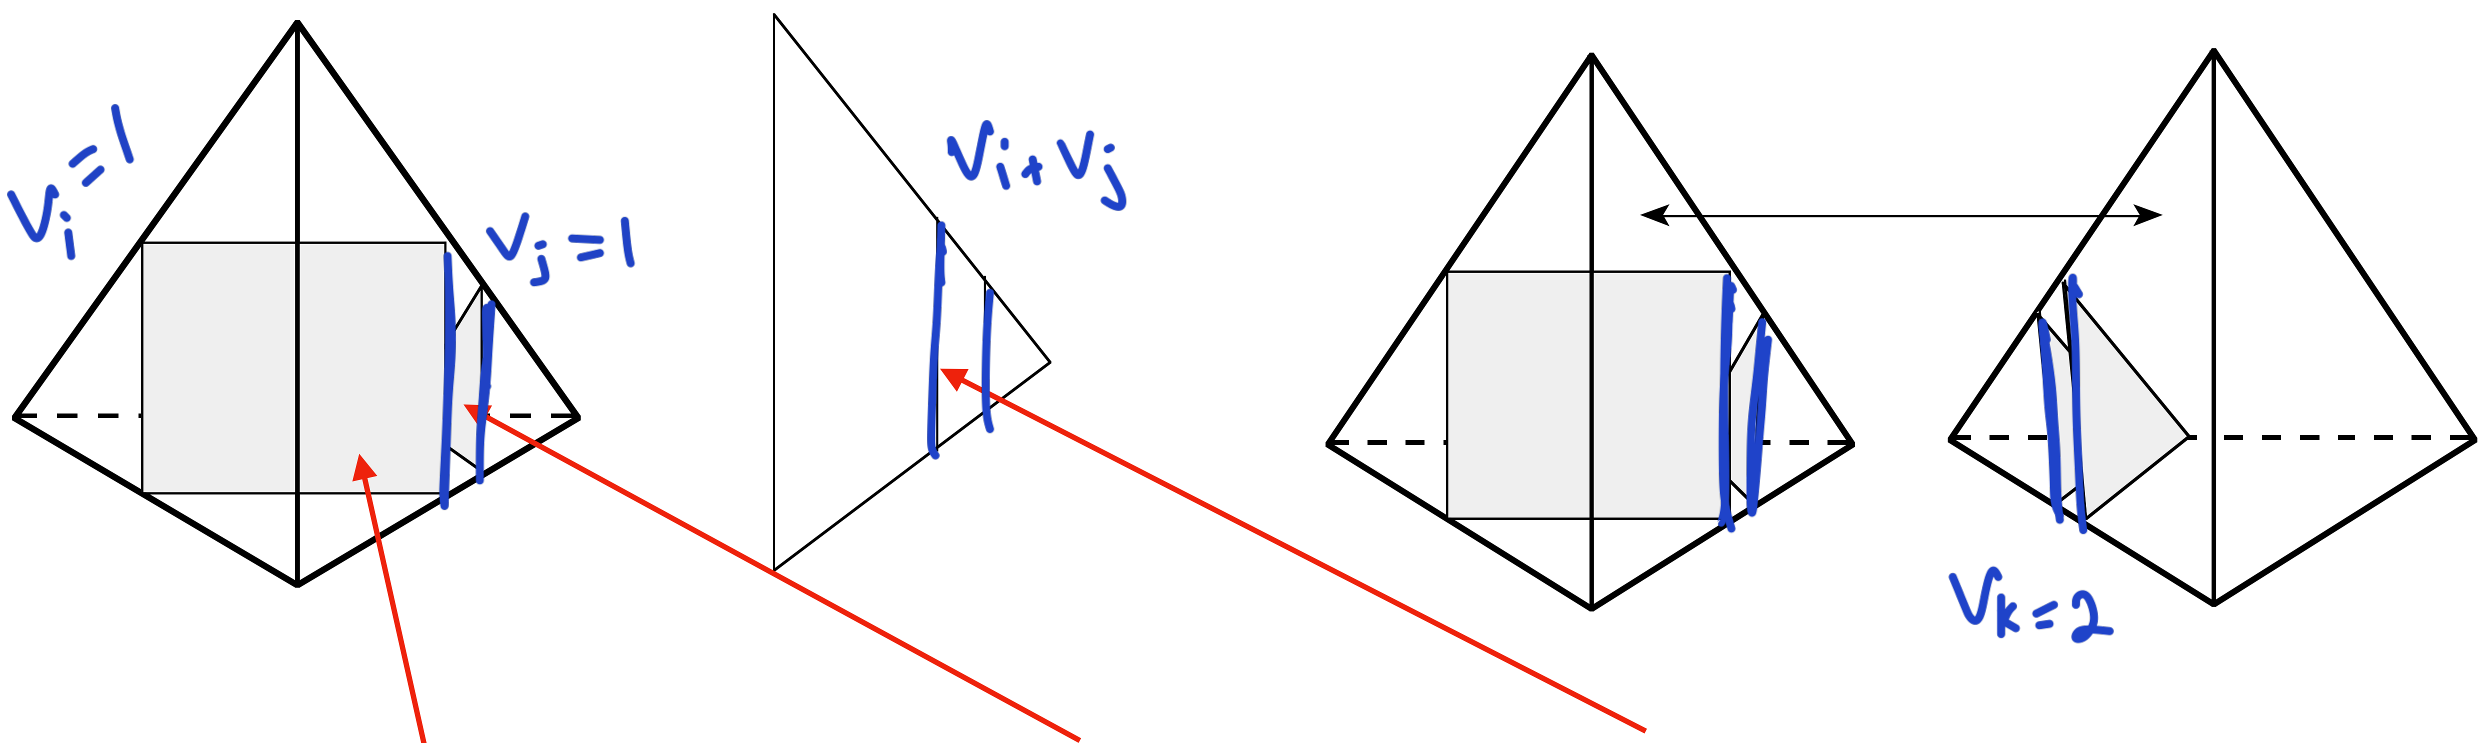

One type of quadrilateral and one type of triangle contribute to edges of this type on this face of a tetrahedron. The number of such edges matches along adjacent tetrahedra.

This again leads to *integer linear programming*.

Matching Equations

Starting with any triangulation having t tetrahedra, we get:

$(v_1, v_2, v_3, \ldots\ v_{7t})$ in $\mathbf{Z}_+^{7t}$ 7t variables

$v_i + v_j = v_k + v_l$ 6t equations

$v_i \geq 0$ 7t inequalities

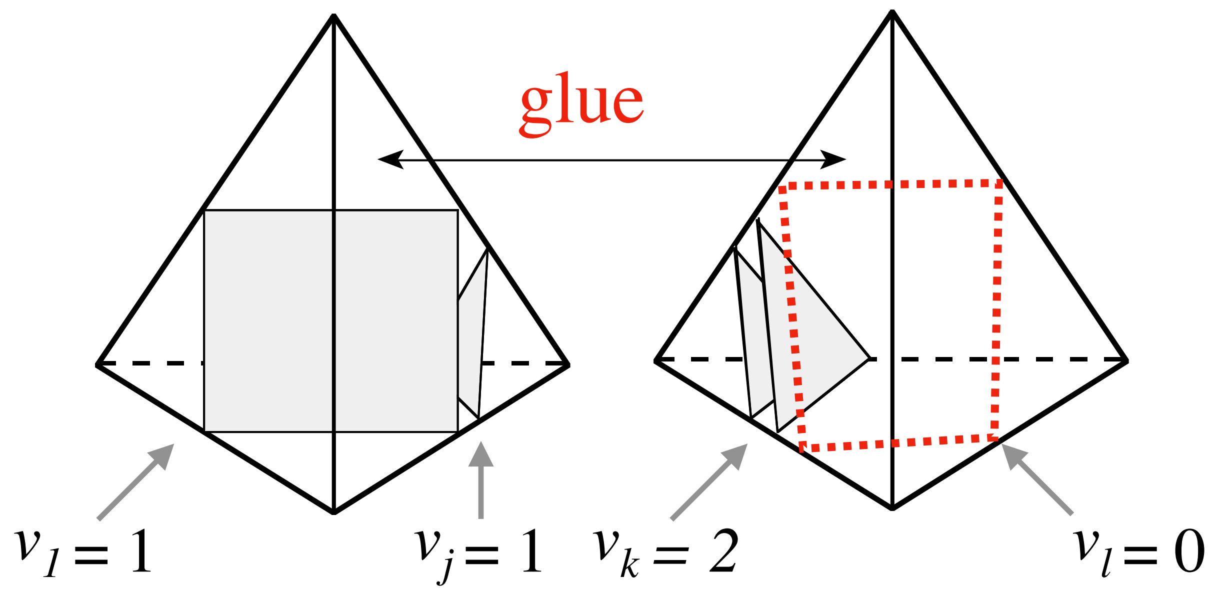

Integer linear programming implies that there are finitely many Fundamental Solutions, and these can be explicitly constructed.

Note that we haven't yet specified a question or algorithm.

Geometric and Algebraic Sums

As with curves, there is a correspondence between

1. Sums of normal vectors (algebra)
2. Sums of normal surfaces (geometry)

called *Haken sum* or *regular sum*

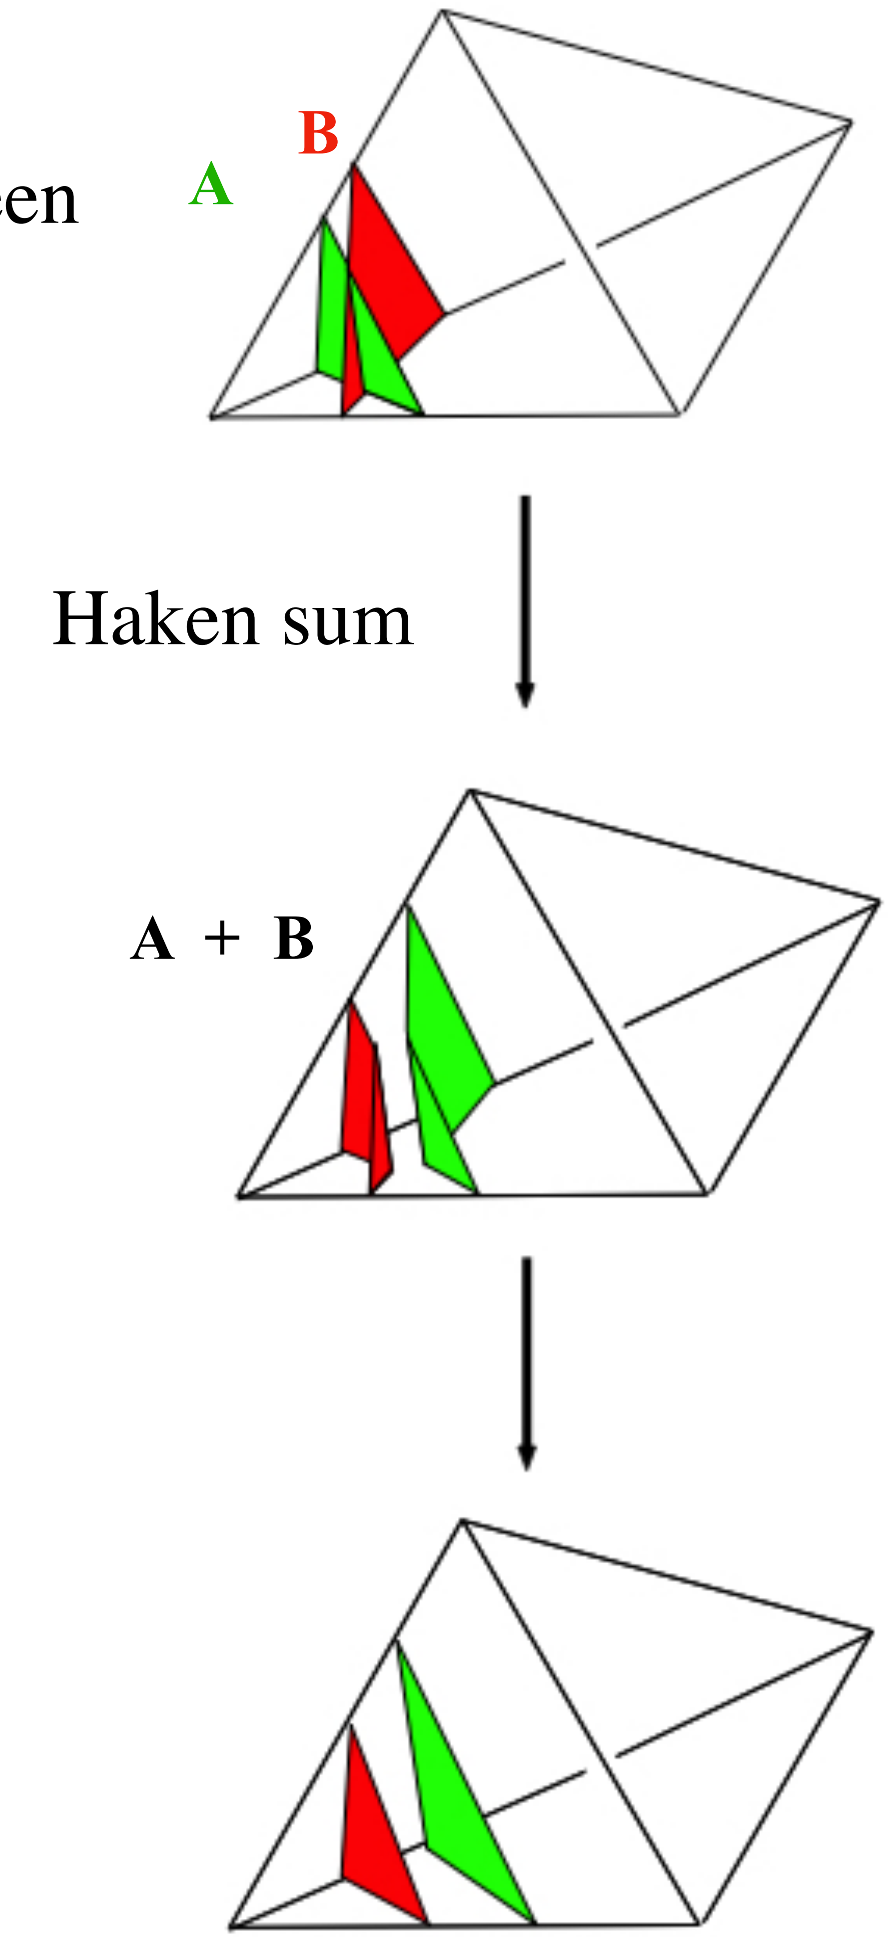

< ...., 0, 1, 0, 0, 0, 0, ... >

\+

< ...., 0, 1, 0, 0, 0, 0, ... >

=

< ...., 0, 2, 0, 0, 0 ,0,... >

Algebraic sums of normal vectors

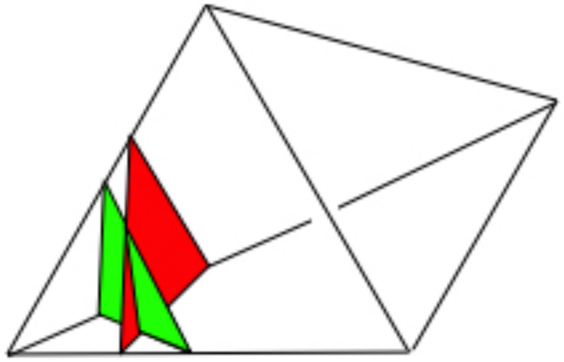

Irregular Sum

Haken Sum

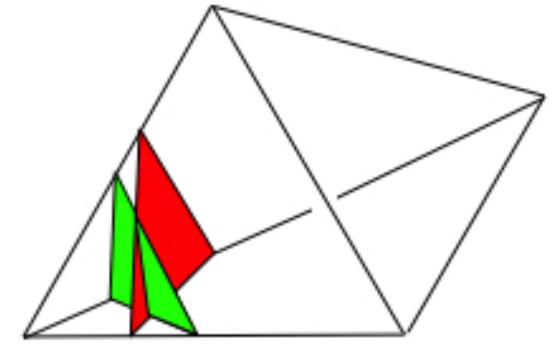

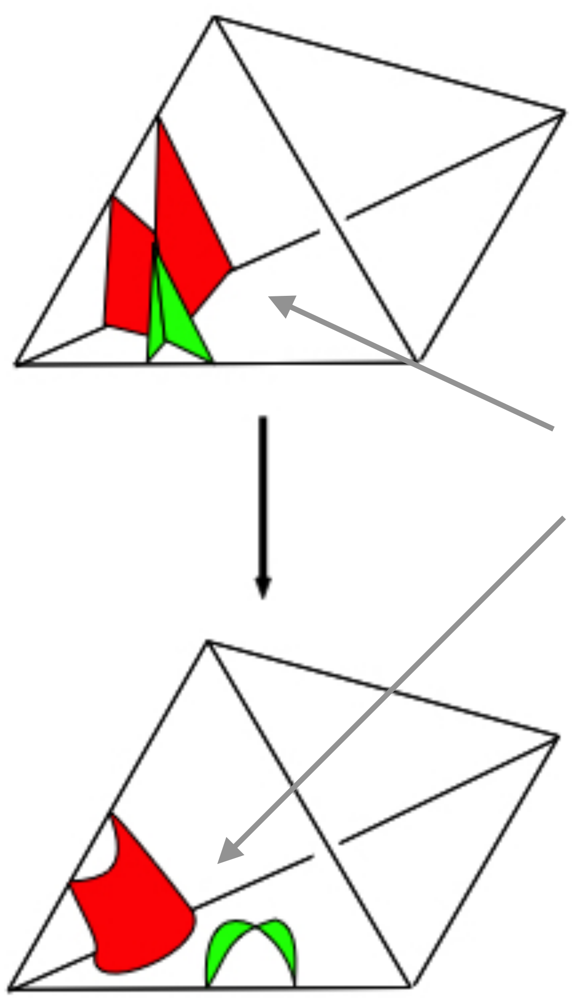

This choice for cut and paste produces a pair of disks that is not normal. We do not want to cut and paste in this way.

In all cases, there is exactly one choice that gives a normal surface.

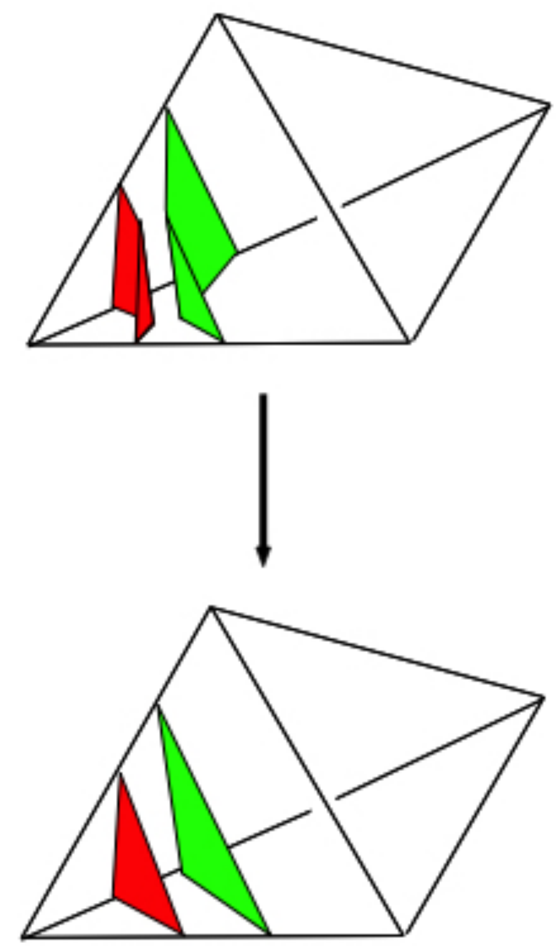

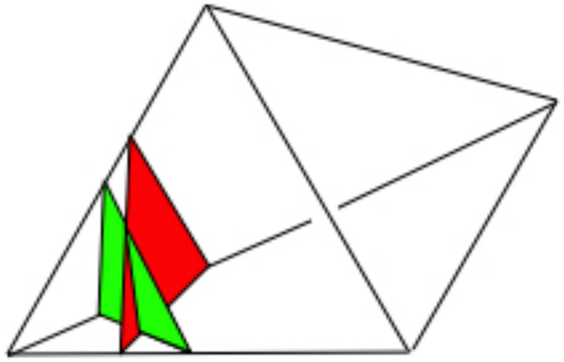

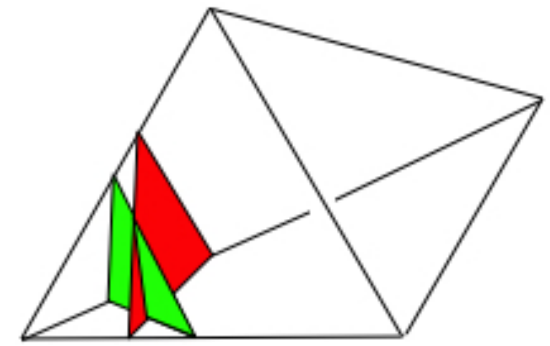

Irregular Sum **Haken Sum**

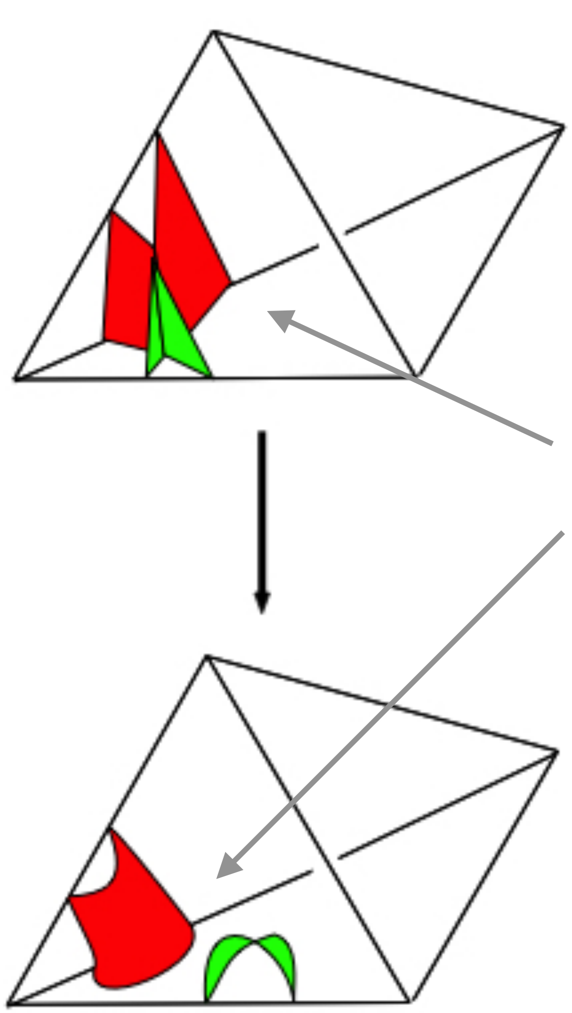

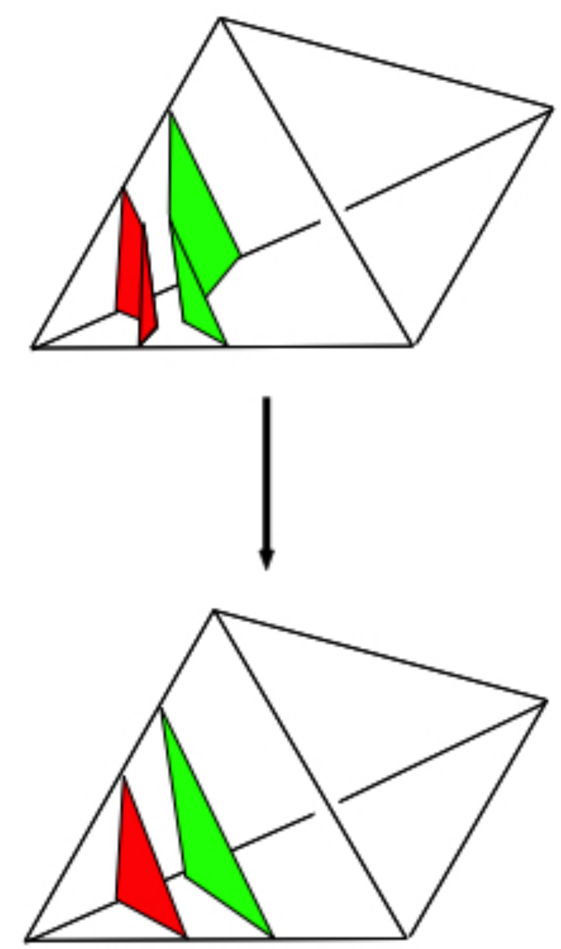

This choice for cut and paste produces a pair of disks that is not normal. We do not want to cut and paste in this way.

In all cases, there is exactly one choice that gives a normal surface.

At this point Saul will ask about quadrilaterals.

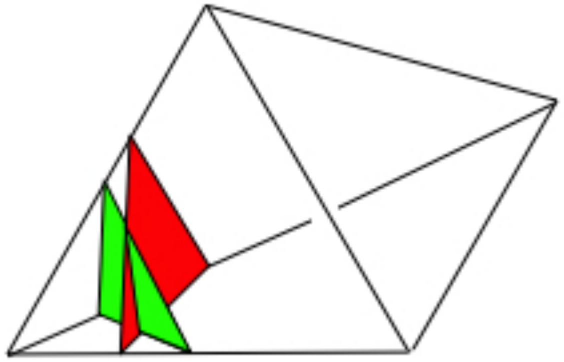

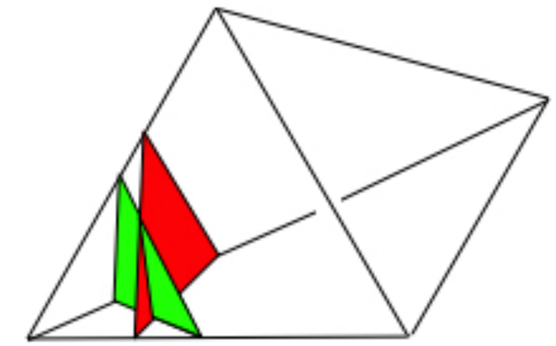

Irregular Sum

Haken Sum

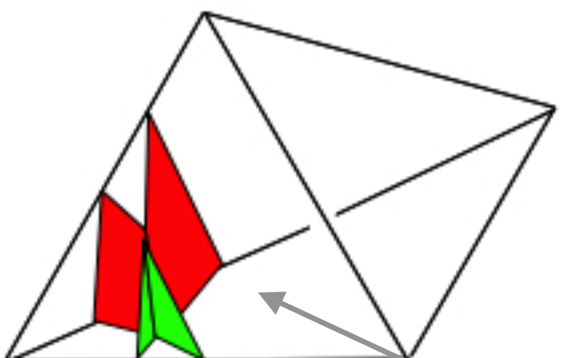

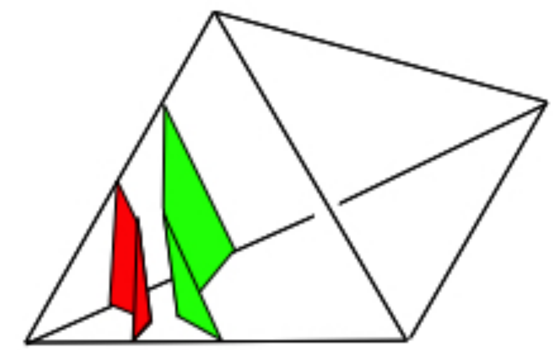

This choice for cut and paste produces a pair of disks that is not normal. We do not want to cut and paste in this way.

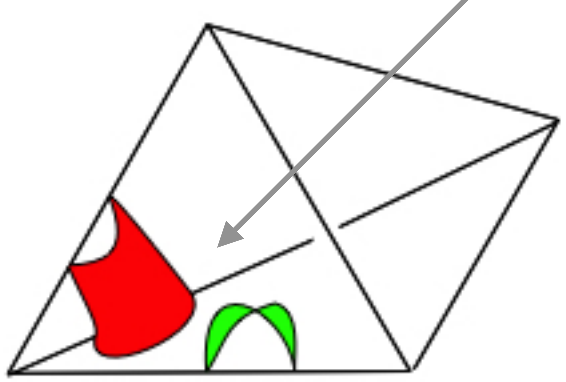

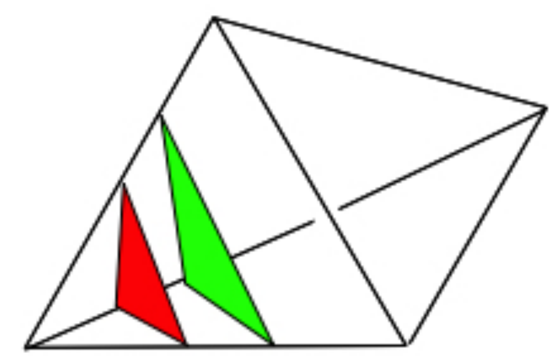

In all cases, there is exactly one choice that gives a normal surface.

At this point Saul will ask about quadrilaterals. Unless he's tired.

Haken (geometric) sums and Euler characteristic

Euler characteristic is additive under the Haken sum of two normal surfaces

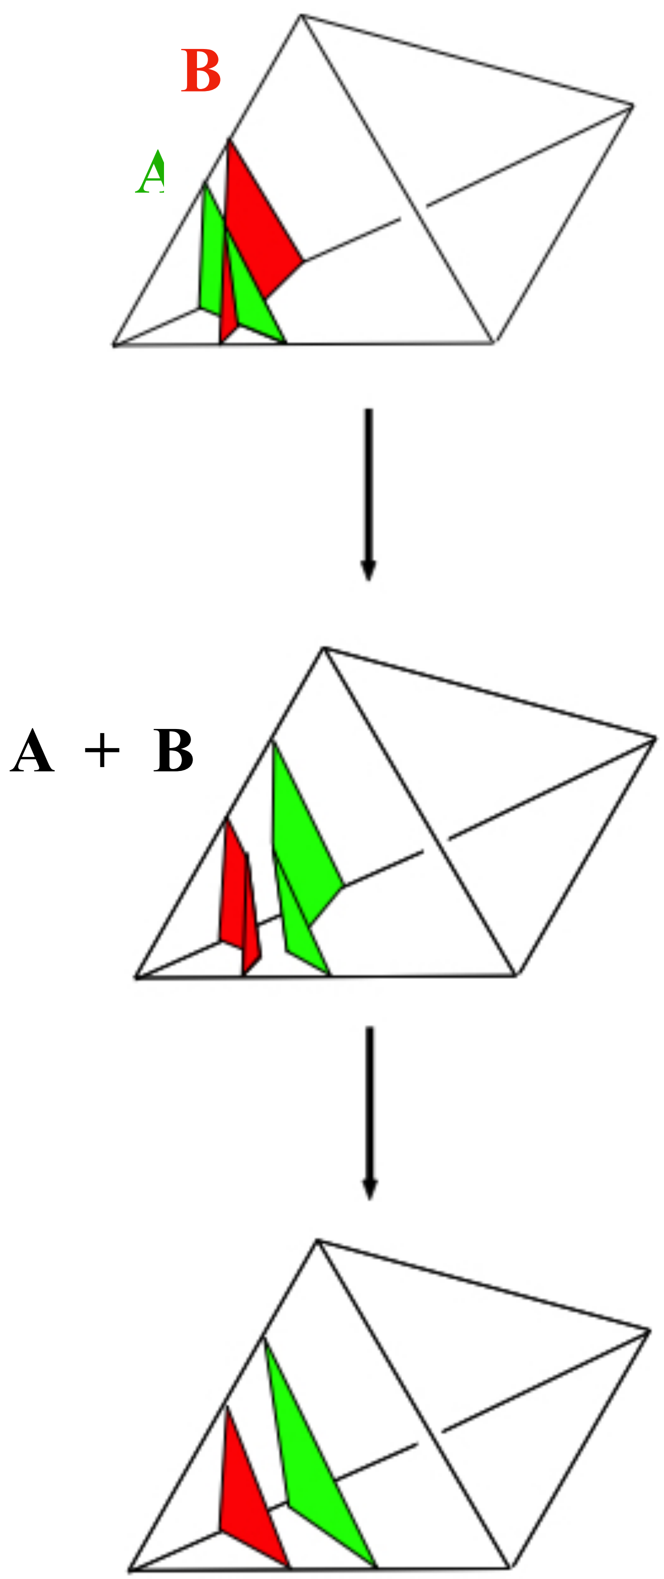

Euler characteristic is computed as V-E+F

If $\mathbf{v} = \mathbf{A} + \mathbf{B}$

then

$$\chi(\mathbf{v}) = \chi(\mathbf{A}) + \chi(\mathbf{B})$$

since the vertices, edges, faces are counted once in each case.

Haken (geometric) sums and weight

Weight is additive under the Haken sum of two normal surfaces

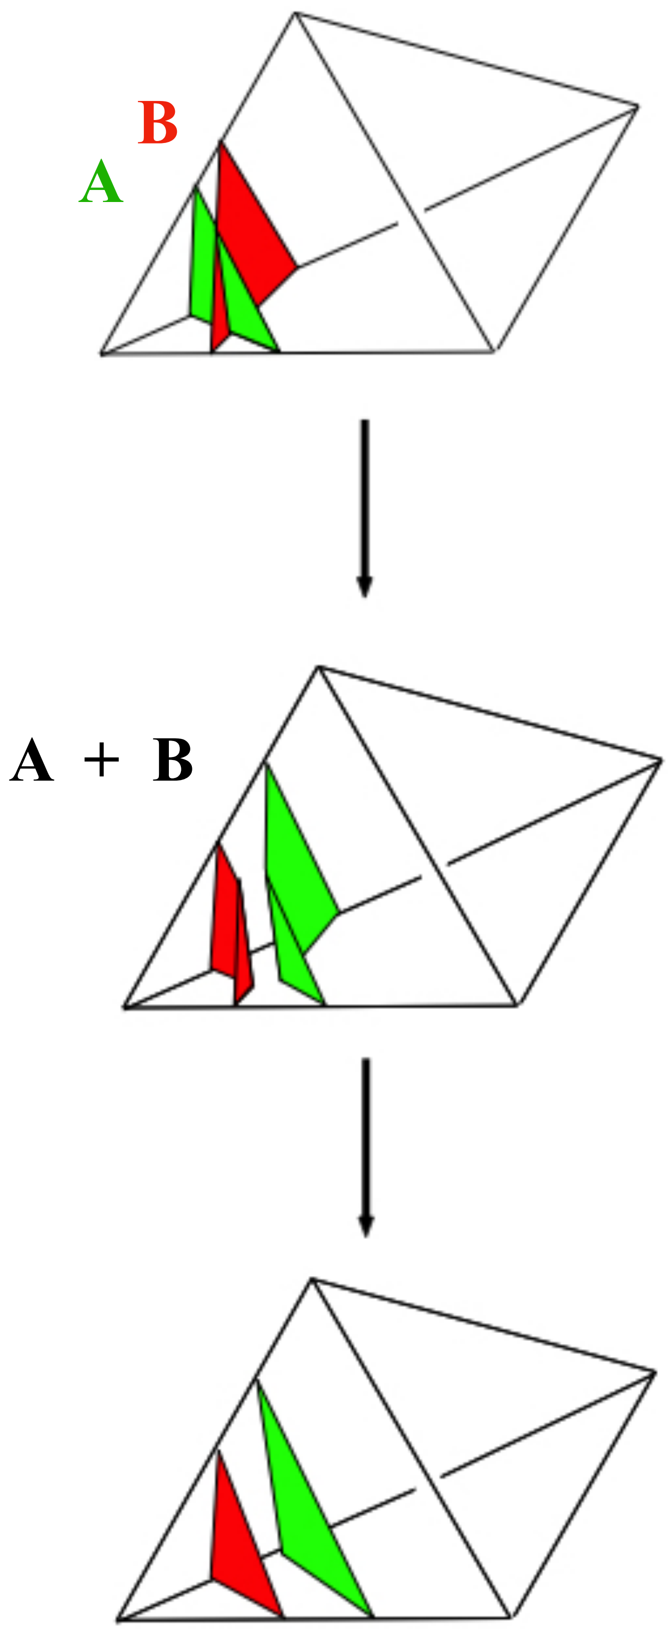

If **v** = **A** + **B**

then

wt(**v**) = wt(**A**) + wt(**B**)

since the number of times that the surfaces cross the 1-skeleton doesn't change.

Reductions of Unknotting

The unknotting problem is a consequence of (reduction from)
Boundary Compressible
Instance: A triangulated 3-manifold with boundary M.
Question: Is the boundary of M compressible?

We will see that Unknotting is also a reduction from the problem
Split Link (Schubert 1961)
Instance: A link L in S^3 with complement M_L.
Question: Does M_L contain a 2-sphere that separates the components of L?

These and many similar problems can be solved with normal surface theory. The algorithms takes the form:

1. Fix a triangulation of M.
2. Construct the Fundamental Normal Surfaces in this triangulation.
3. Check if one of these has the sought after property.

Justification of the algorithm requires a proof that the Fundamental Normal Surfaces contain one of the sought after surfaces, if it exists.

Boundary Conditions for Unknotting

When we look for a normal disk whose boundary is a knot, we need to avoid trivial disks that don't span the knot.

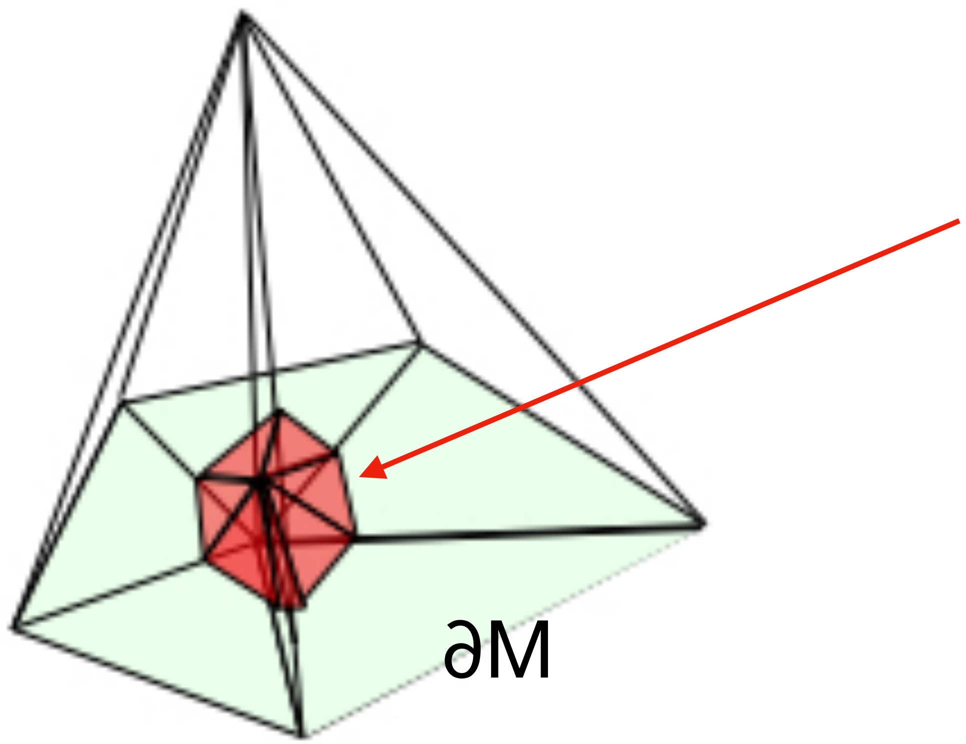

Trivial Normal Disk.
(A hemisphere linking a boundary vertex.)

We are searching for disks whose boundaries run once over the knot. These disks have boundary a longitude on the torus boundary of $M_K = S^3$-K.

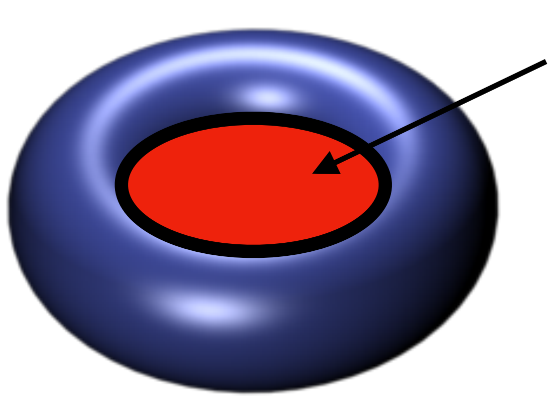

compressing disk

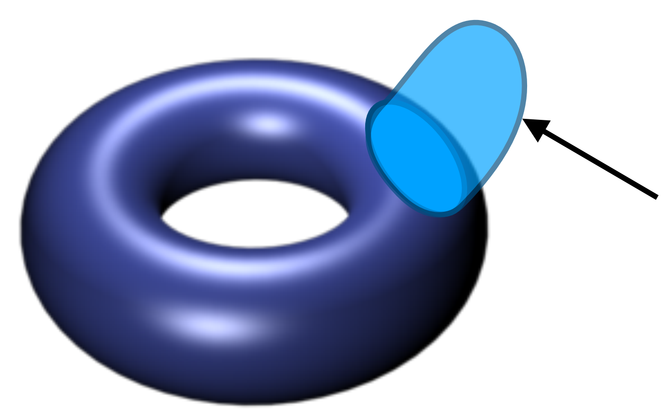

trivial disk

Boundary Conditions for Unknotting

We can ensure that the boundary of the normal vectors we solve for run once around a longitude by setting certain variables equal to zero, namely all triangles and quadrilaterals meeting the boundary except for a collection that runs around a longitude (green)

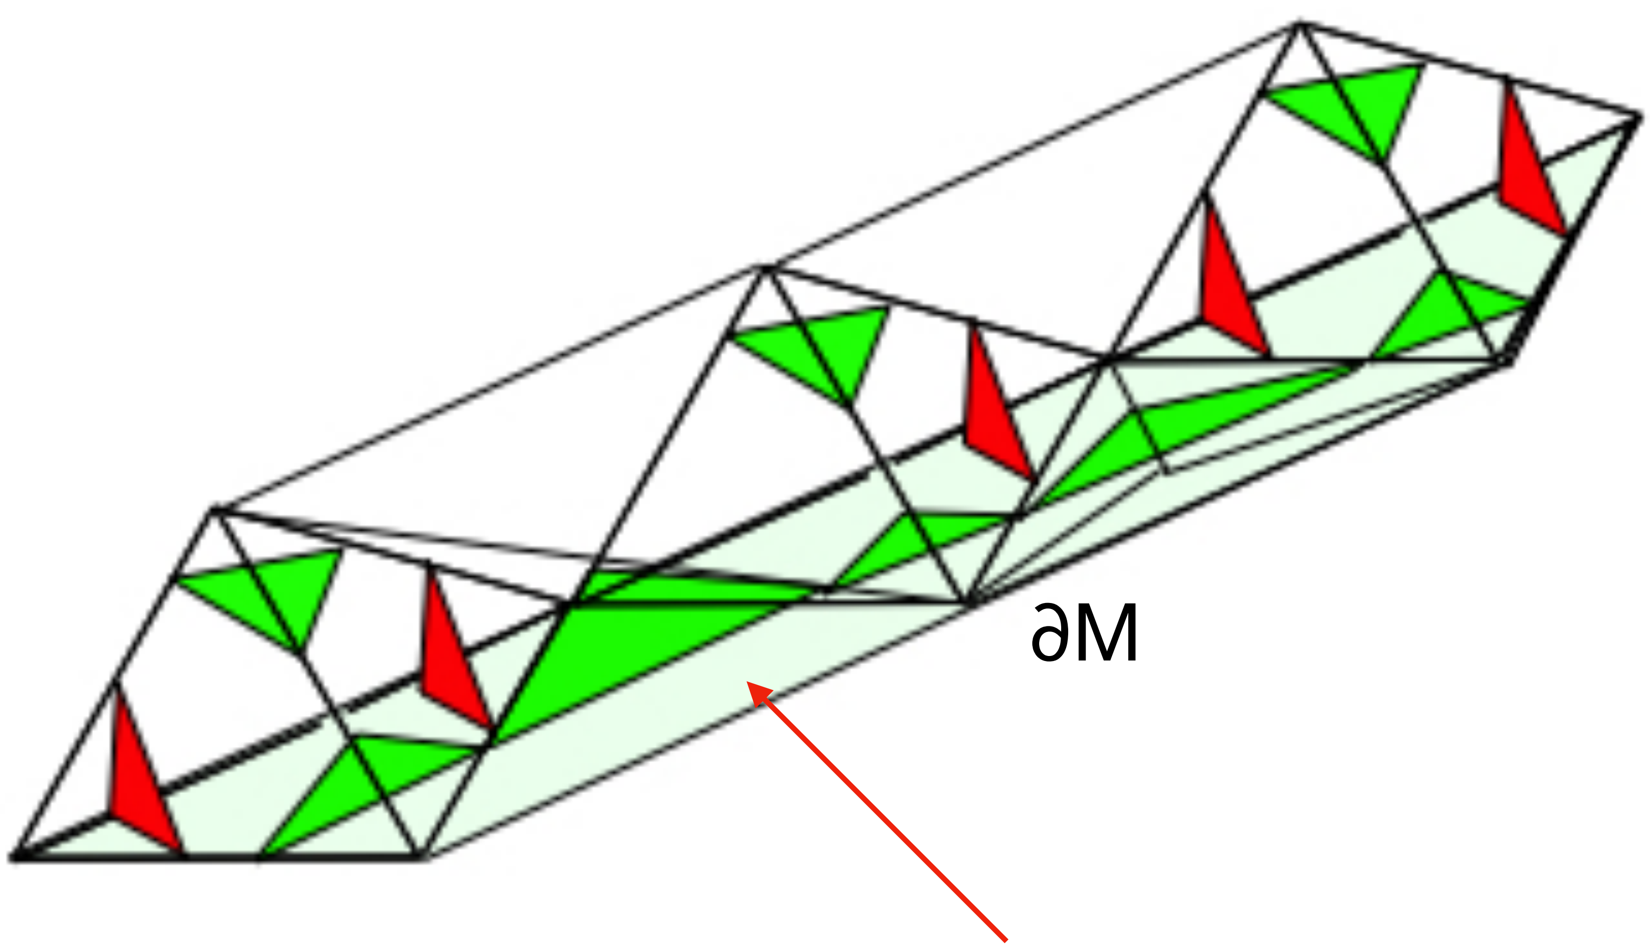

The blue annulus on ∂M contains a longitude

Boundary Conditions for Unknotting

When we look for a disk whose boundary is a knot, we can avoid generating boundary parallel disks by setting some variables $v_i = 0$. We can set all variables corresponding to triangles or quadrilaterals that meet ∂M to zero (red), except for those along a longitude, which we allow to have any value.

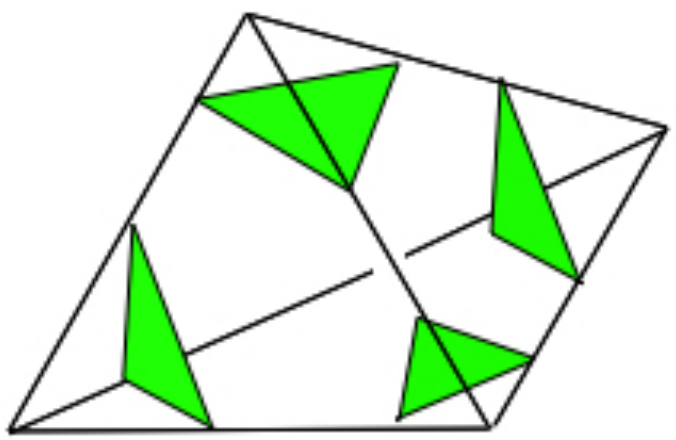

Interior Tetrahedron

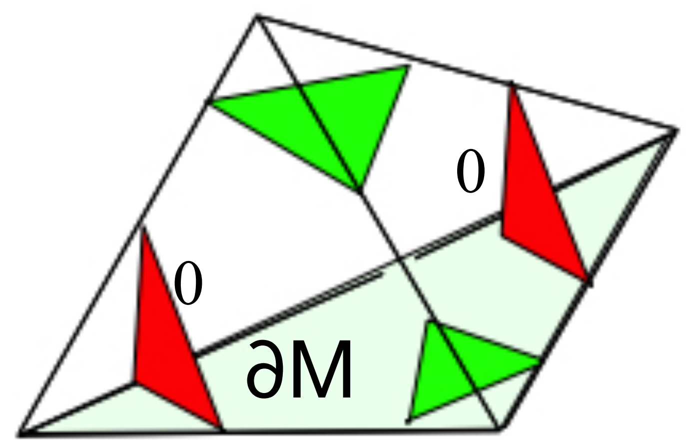

Tetrahedron meeting ∂M on a face

Boundary Conditions for Unknotting

When we look for a disk whose boundary is a knot, we can avoid generating boundary parallel disks by setting some variables $v_i = 0$. We can set all variables corresponding to triangles or quadrilaterals that meet ∂M to zero (red), except for those along a longitude, which we allow to have any value.

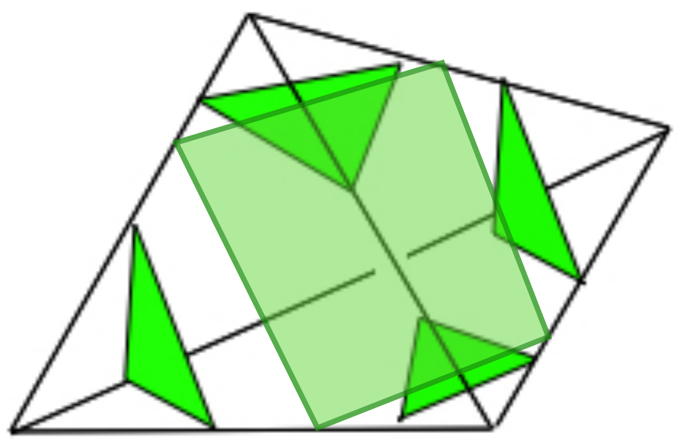

Interior Tetrahedron

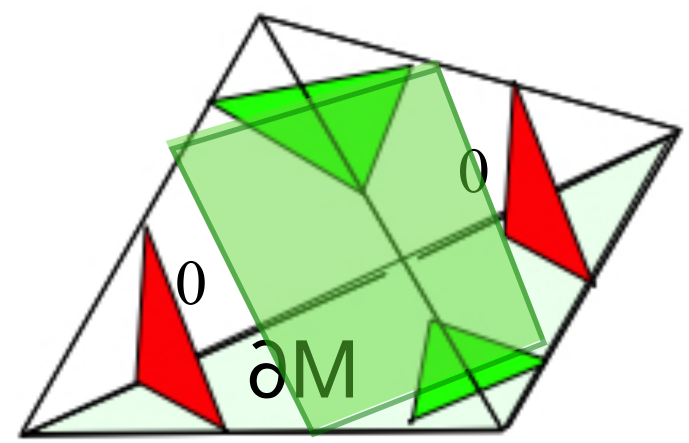

Tetrahedron meeting ∂M on a face

We may need to subdivide the triangulation to realize a longitude this way.

Alternately, we can avoid worrying about boundary conditions by reformulating the unknotting problem as a problem about closed surfaces:

SPLIT LINK

Problem: SPLIT LINK
Instance: A link L in S^3 with complement M_L.
Question: Does M_L contain a 2-sphere that separates the components of L?

Claim. Unknotting for K reduces to Split Link for L

Given a knot K, create a 2-component link L by taking a pushed-off copy of K (chosen so that the linking number is zero).

K is unknotted if and only if there is a 2-sphere in S^3-L separating the two components of L.

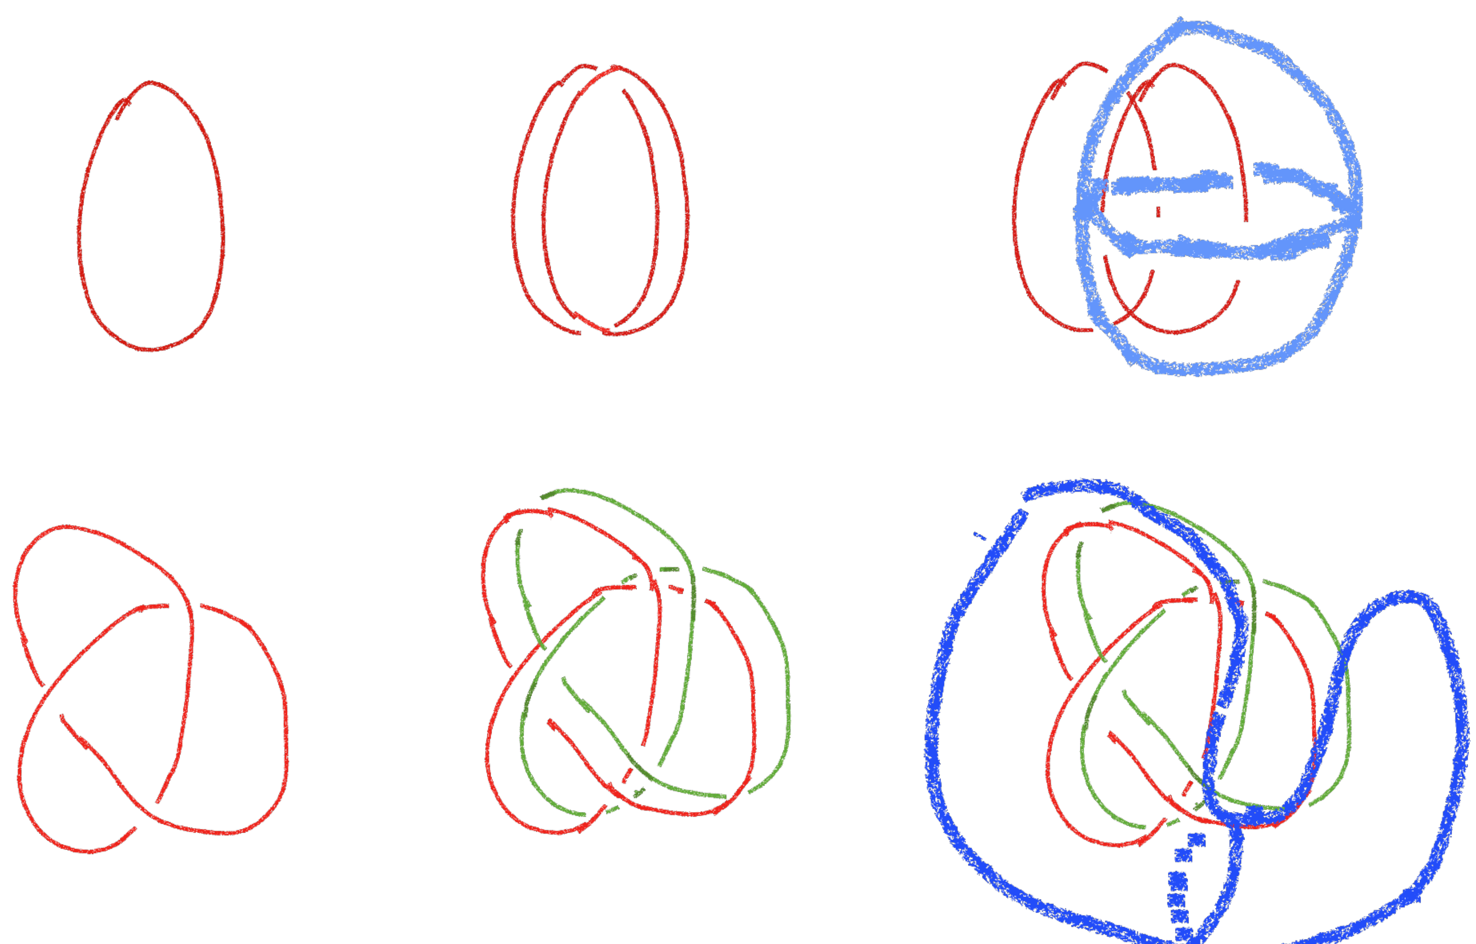

Can this link be split?

SPLIT LINK

Problem: SPLIT LINK
Instance: A link L in S^3-L with complement M_L.
Question: Does M_L contain a 2-sphere that separates the components of L?

Claim. Unknotting for K reduces to Split Link for S^3-L

Given a knot K, create a 2-component link L by taking a pushed-off copy of K (chosen so that the linking number is zero).

If there is a 2-sphere in S^3-L separating the two components of L, the 2-sphere intersects the annulus between the two components of L in an essential curve. This curve is isotopic to K and bounds a disk. So K is unknotted.

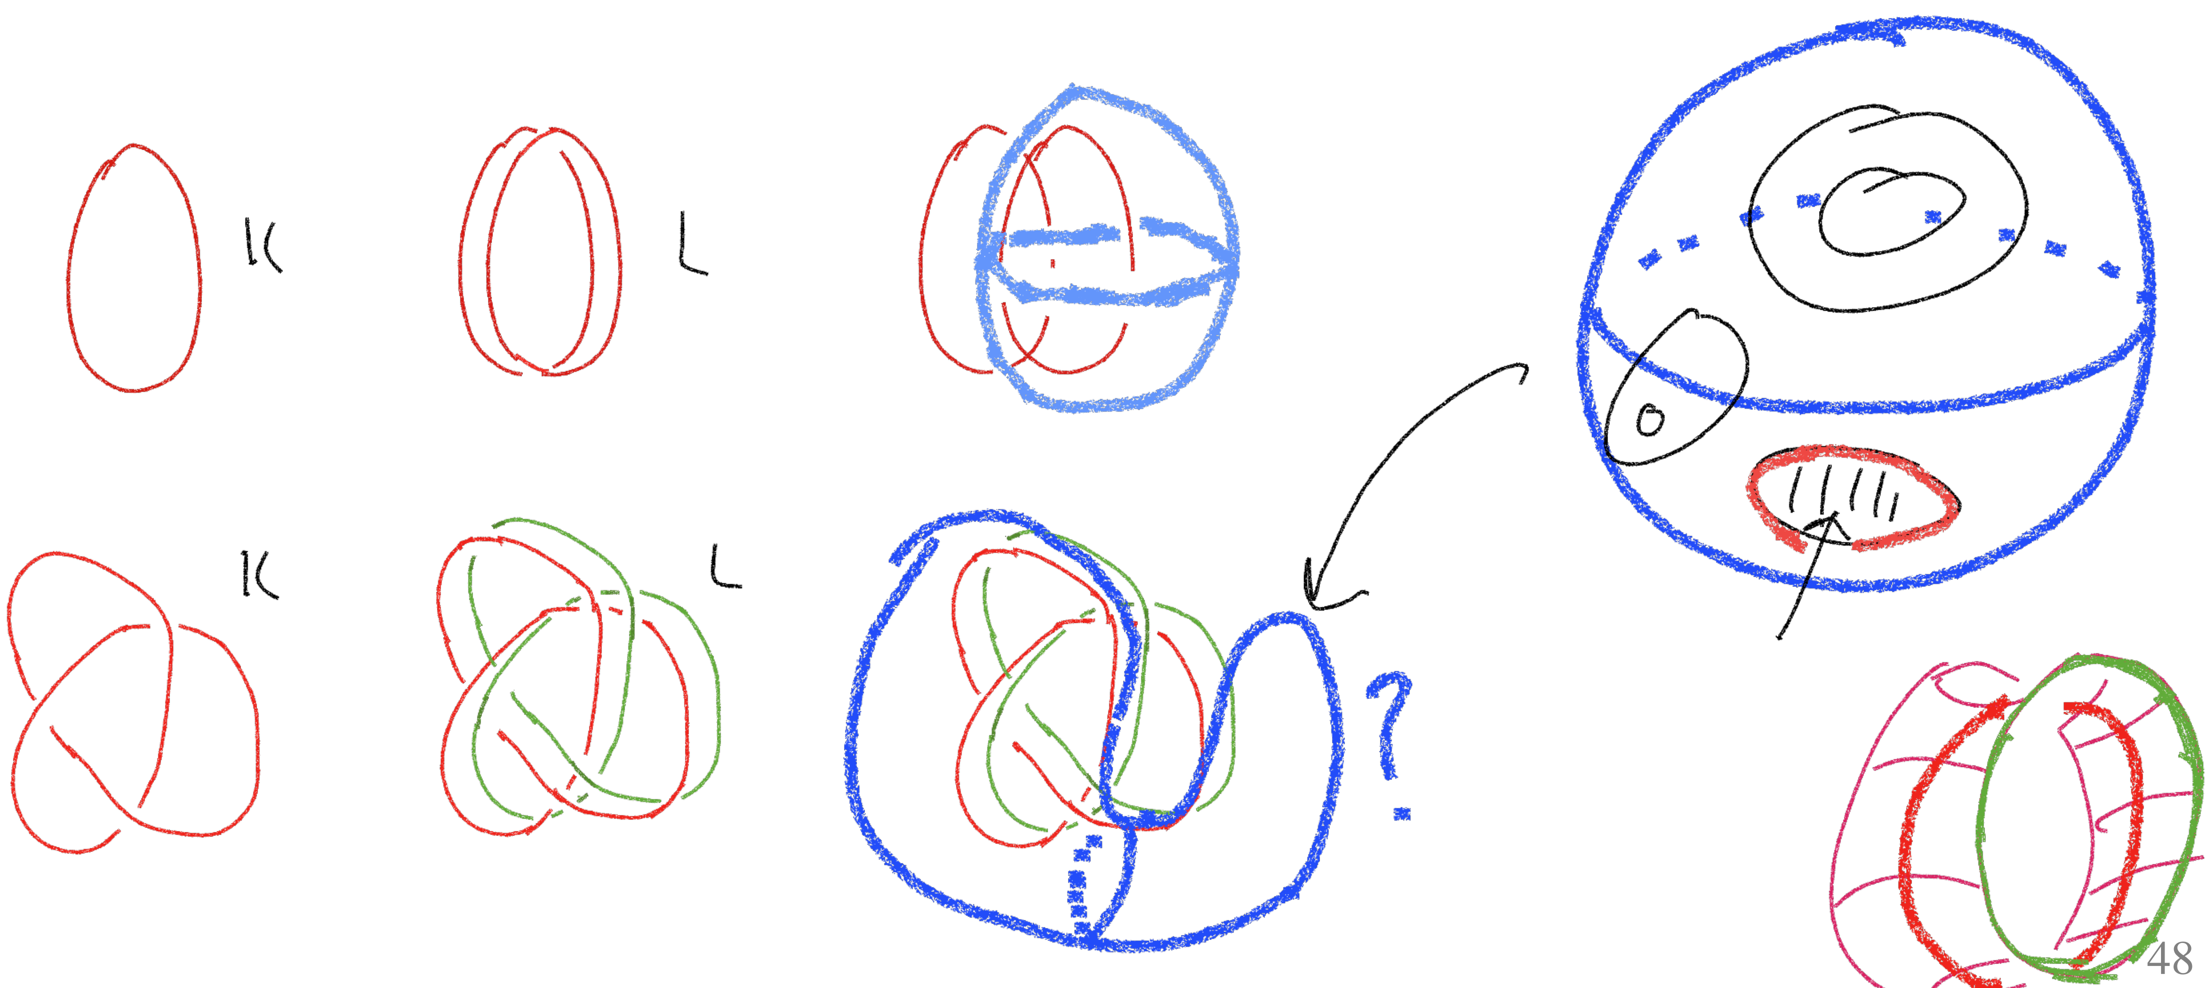

An Algorithm for SPLIT LINK

SPLIT LINK
Instance: A link L in S^3 with complement M_L.
Question: Does M_L contain a 2-sphere that separates the components of L?

The Algorithm

1. Construct all fundamental solutions to the normal surface equations.
2. Check whether any of the fundamental solutions has $\chi(\mathbf{v}) = 2$ and separates components of L. If one does, then answer “YES” Otherwise answer “NO”.

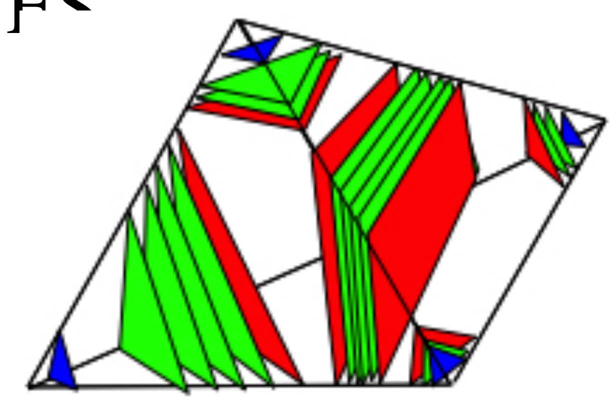

Proof that the algorithm works as claimed:
Need to show that if there is a splitting 2-sphere then there is a fundamental vector **v** that represents a splitting 2-sphere. There are finitely many fundamental solutions and each can be checked to see if it is a 2-sphere that separates components of L.

Corollary: UNKNOTTING

Fundamental Surfaces

Lemma: If L is split then M_L contains a fundamental normal splitting 2-sphere.

A fundamental normal surface has associated vector **v** that cannot be written as a sum of two non-zero normal vectors. $\mathbf{v} \neq \mathbf{A} + \mathbf{B}$.

Proof: If L is split then M_L contains a splitting 2-sphere.

Step 1. **Normalize**: If L is split then $M_L = S^3 - L$ contains a *normal* splitting 2-sphere.

The normalization procedure shows that any surface in a triangulated 3-manifold can be transformed to a normal surface by a sequence of the following moves:

1. Isotopy in M_L
2. Compression and boundary-compression in M_L
3. Eliminating components lying inside a single tetrahedron.

Start with a splitting 2-sphere S. Isotopy moves it around, but it still separates components of L. Boundary compression only applies to surfaces with boundary.

If S compresses into a pair of 2-spheres, then at least one of these is also a splitting 2-sphere.

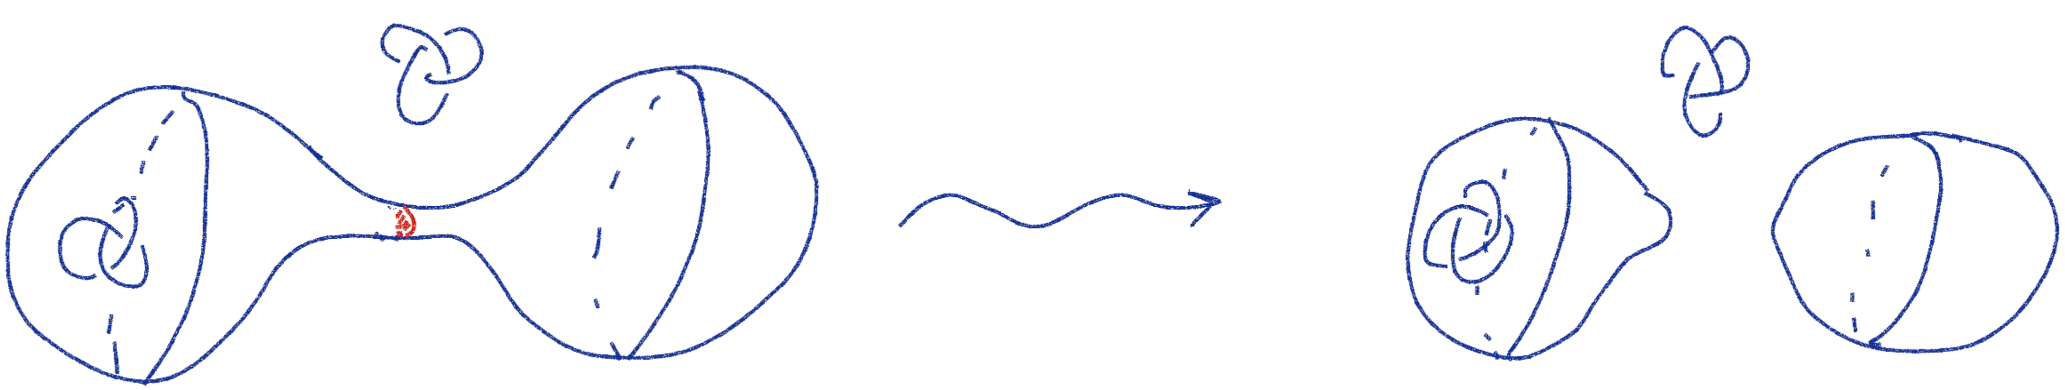

So normalization yields a normal splitting 2-sphere in M_L.

Step 2. Fundamentalize

Lemma: If L is split then M_L contains a fundamental normal splitting 2-sphere.

Step 2. **Fundamentalize**: If M_L contains a normal splitting 2-sphere then it contains a *fundamental* normal splitting 2-sphere.

Choose a normal splitting 2-sphere S with normal vector $\mathbf{v} = (v_1, v_2, v_3, \ldots\ v_{7t})$, of *smallest weight* among all normal splitting 2-spheres. This 2-sphere is fundamental.

Suppose **v** is not fundamental. Then $\mathbf{v} = \mathbf{A} + \mathbf{B}$ for some normal vectors **A** and **B**.

Lemma (Schubert):
Let **v** represent a connected normal surface and suppose that **A** and **B** represent two normal surfaces and **v** is the Haken sum of **A** and **B** along the collection of curves $\mathbf{A} \cap \mathbf{B}$, $\mathbf{v} = \mathbf{A} + \mathbf{B}$.
If **A** and **B** are chosen to minimize $|\mathbf{A} \cap \mathbf{B}|$ then

i. **A** and **B** each represent connected surfaces and
ii. No curve of intersection of $\mathbf{A} \cap \mathbf{B}$ is separating on both **A** and on **B.**

Schubert's Lemma

Lemma (Schubert 1961):
Let **v** be a normal vector representing a connected normal surface. Suppose that **A** and **B** represent two normal surfaces and **v** is the Haken sum of **A** and **B**, **v** = **A** + **B**.
If **A** and **B** are chosen to minimize $|A \cap B|$ then

i. **A** and **B** each represent connected surfaces and
ii. No curve of intersection of $A \cap B$ is separating on both **A** and on **B.**

Proof
i. Suppose not. Then we can write $\mathbf{B} = \mathbf{B}_1 + \mathbf{B}_2 + \mathbf{B}_3 + \ldots \mathbf{B}_k$ where each $\mathbf{B}_i$ is a connected normal surface, $k>2$. Take Haken sum along intersection curves of $\cup \mathbf{B}_i$ until obtaining exactly two intersecting embedded normal surfaces **C** and **D**.

Then $\mathbf{v} = \mathbf{C} + \mathbf{D}$ and $|\mathbf{C} \cap \mathbf{D}| < |\mathbf{A} \cap \mathbf{B}|$, a contradiction.

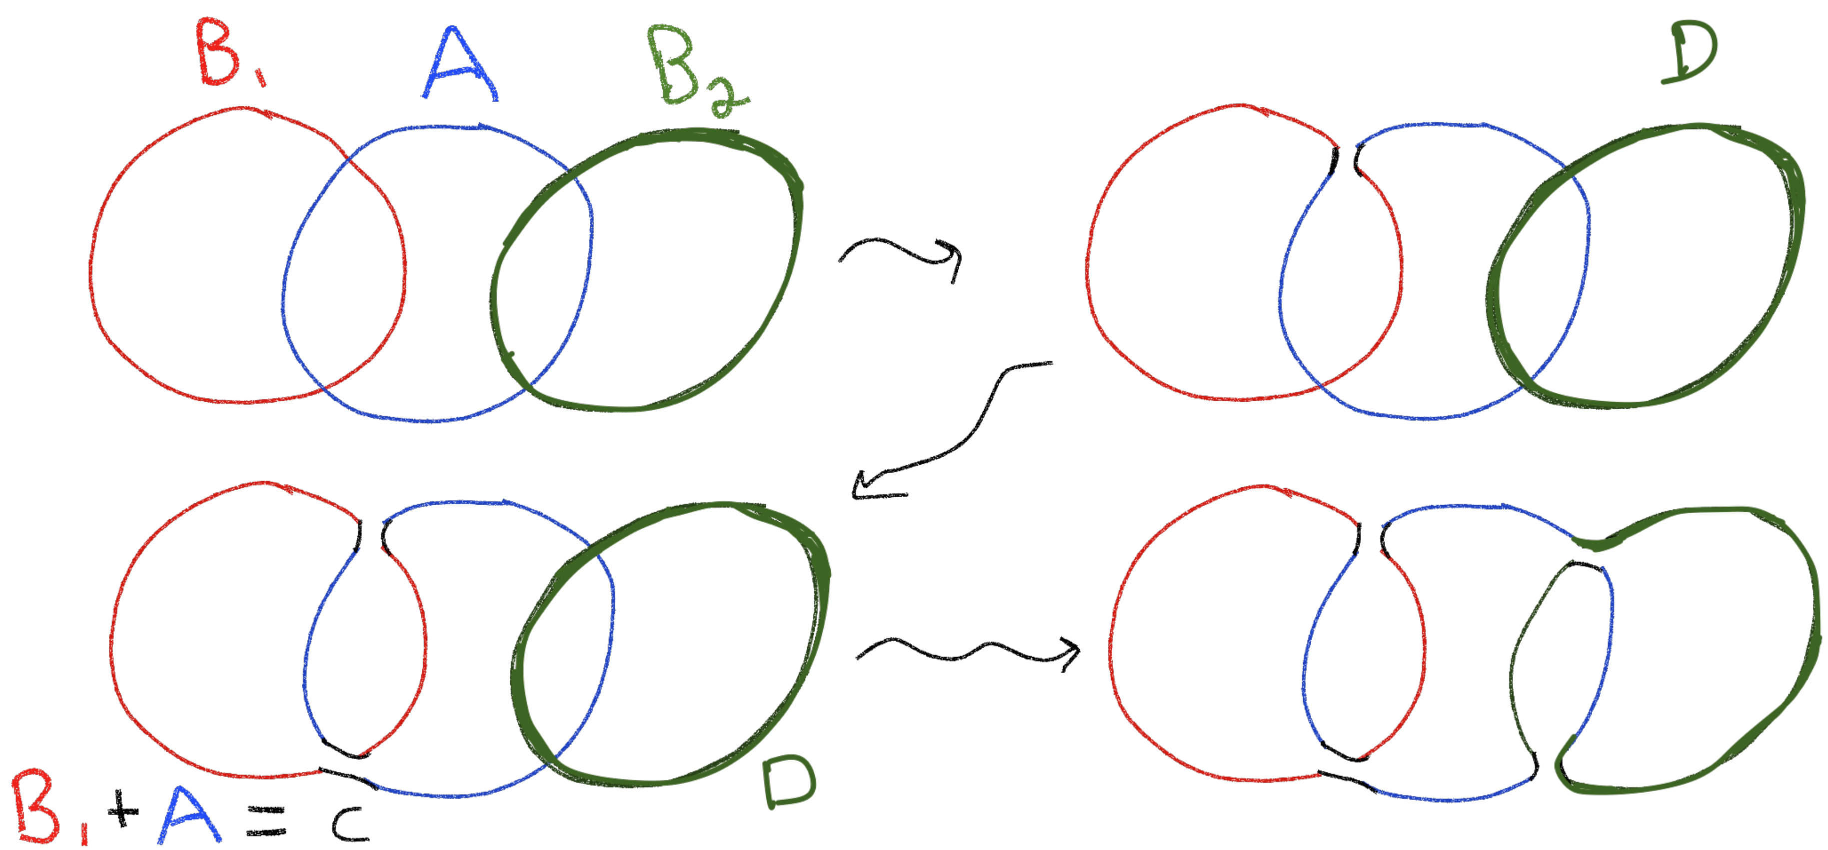

Haken Sum along all curves gives us **v**.

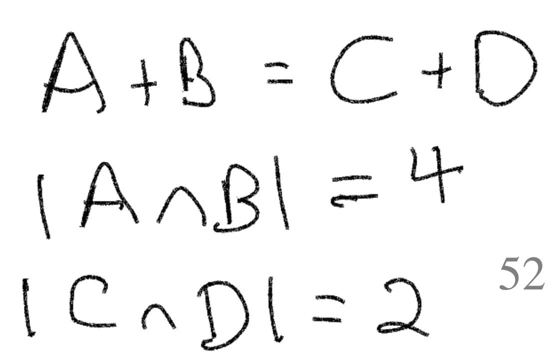

Schubert's Lemma

Lemma (Schubert):
Let **v** be a normal vector representing a connected normal surface. Suppose that **A** and **B** represent two normal surfaces and **v** is the Haken sum of **A** and **B**, **v** = **A** + **B**.
If **A** and **B** are chosen to minimize |**A**∩**B**| then

i. **A** and **B** each represent connected surfaces and
ii. No curve of intersection of **A**∩**B** is separating on both **A** and on **B.**

Proof
i. Suppose not. Then we can write $\mathbf{B} = \mathbf{B}_1 + \mathbf{B}_2 + \mathbf{B}_3 + \ldots \mathbf{B}_k$ where each $\mathbf{B}_i$ is a connected normal surface, $k>2$. Take Haken sum along intersection curves of $\cup \mathbf{B}_i$ until obtaining exactly two intersecting embedded normal surfaces **C** and **D**.

Then **v** = **C** + **D** and |**C**∩**D**| < |**A**∩**B**|, a contradiction.

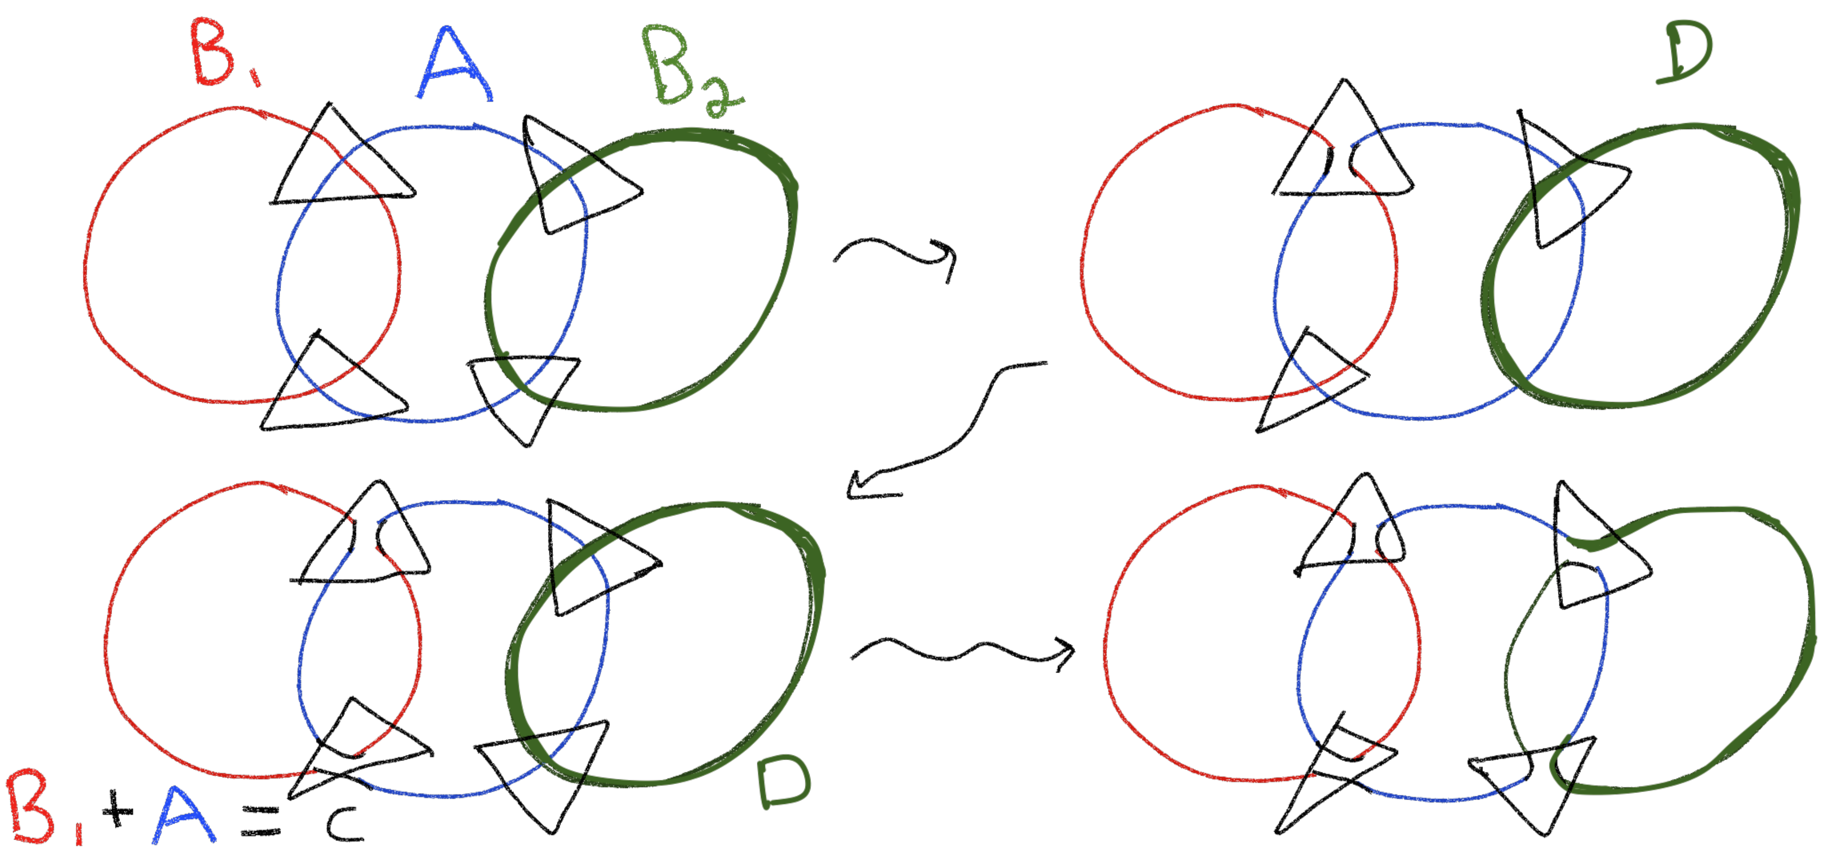

Haken Sum along all curves gives us **v**.

A + B = C + D
|A ∩ B| = 4
|C ∩ D| = 2

Schubert's Lemma

Lemma (Schubert):
Let **v** represent a connected normal surface and suppose that **A** and **B** represent two normal surfaces and **v** is the Haken sum of **A** and **B** along the collection of curves **A**∩**B**, **v** = **A** + **B**.
If **A** and **B** are chosen to minimize |**A**∩**B**| then

i. **A** and **B** each represent connected surfaces and
ii. No curve of intersection of **A**∩**B** is separating on both **A** and on **B.**

Proof
ii. If a curve of **A**∩**B** separates on both, then take the Haken sum along this curve. This results in a normal surface with two "components", possibly immersed. Take further Haken sum along intersection curves for as long as it is possible to do so without getting a connected surface.

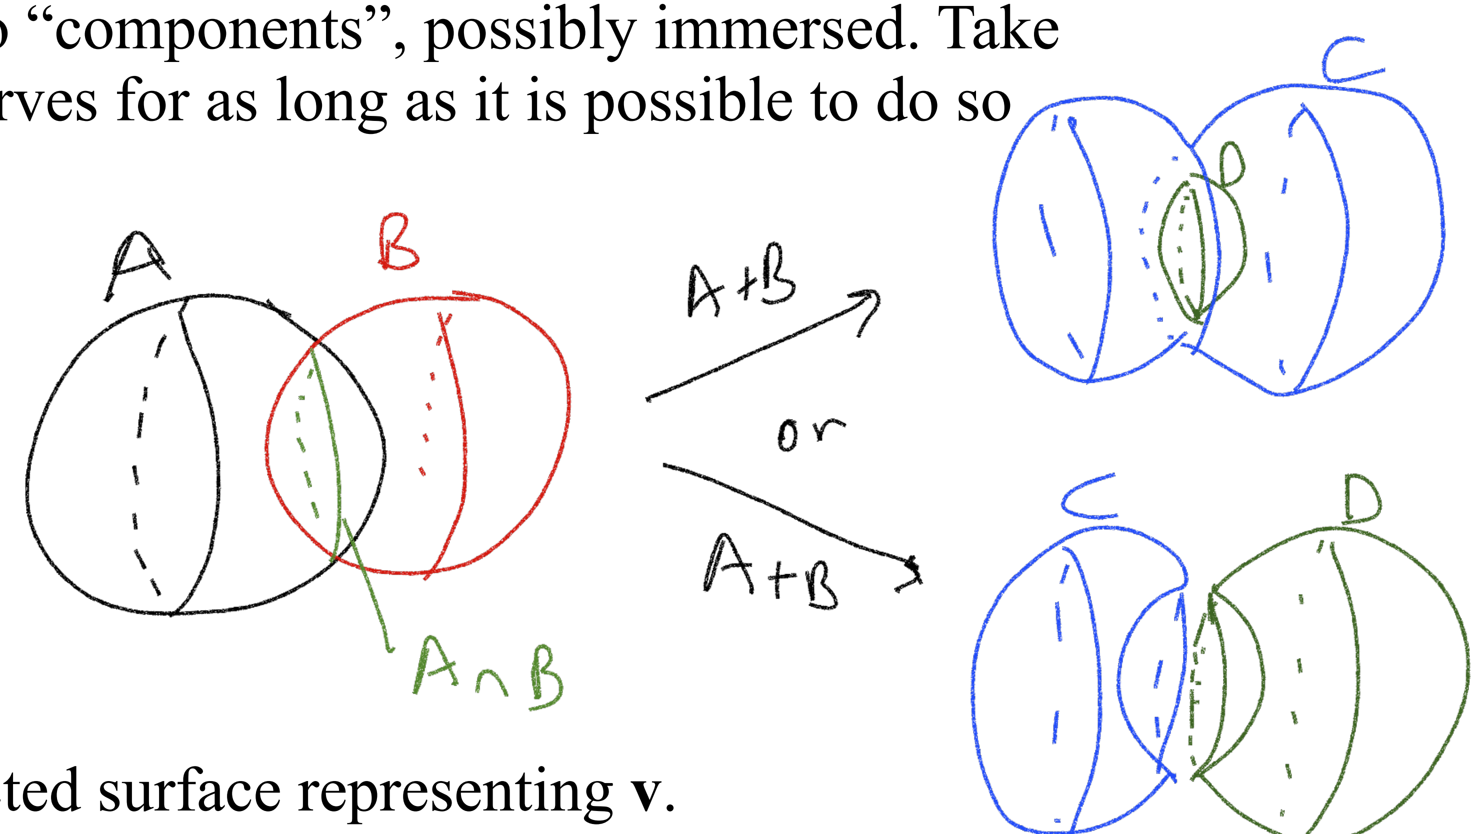

Doing *all* Haken sums result in a connected surface representing **v**.
Each Haken sum changes the number of "components" by at most one.
So Haken sums along some subset of all the intersection curves results in exactly two embedded normal surfaces **C** and **D**: **v** = **C** + **D**
But |**C**∩**D**| < |**A**∩**B**|, a contradiction.

Fundamentalization

Lemma: If L is split then M_L contains a fundamental normal splitting 2-sphere.

Step 2. **Fundamentalize**: If M_L contains a normal splitting 2-sphere then it contains a *fundamental* normal splitting 2-sphere.

Choose a normal splitting 2-sphere S with normal vector $\mathbf{v} = (v_1, v_2, v_3, \ldots\ v_{7t})$, of **smallest weight** among all normal splitting 2-spheres.
Suppose **v** is not fundamental. Then $\mathbf{v} = \mathbf{A} + \mathbf{B}$ for some normal vectors **A** and **B**.
We can assume that each is connected by choosing to minimize $|\mathbf{A} \cap \mathbf{B}|$.
Since Euler characteristic adds under geometric sum,

$$\chi(\mathbf{v}) = \chi(\mathbf{A}) + \chi(\mathbf{B})$$

$$2 = \chi(\mathbf{A}) + \chi(\mathbf{B})$$

So one of **A** and **B** represents a normal 2-sphere (say **A**) and the other a normal torus.
Since each of **A** and **B** has smaller weight than **v**, **A** must represent a non-splitting sphere.

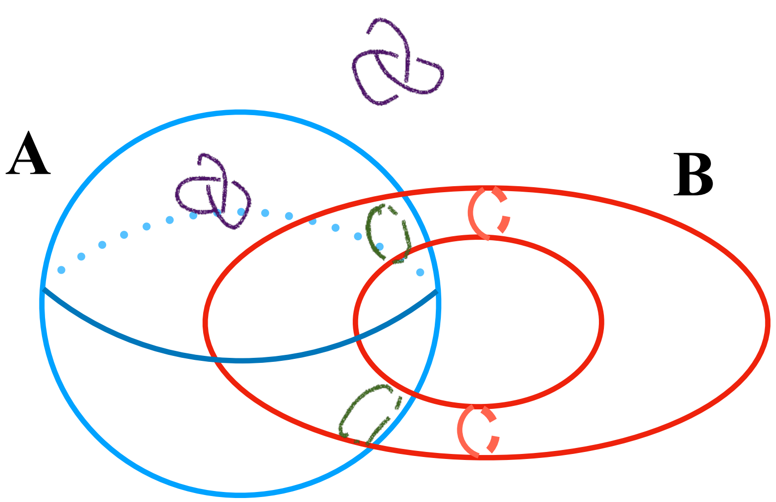

A has smaller weight than **v**
Can't happen

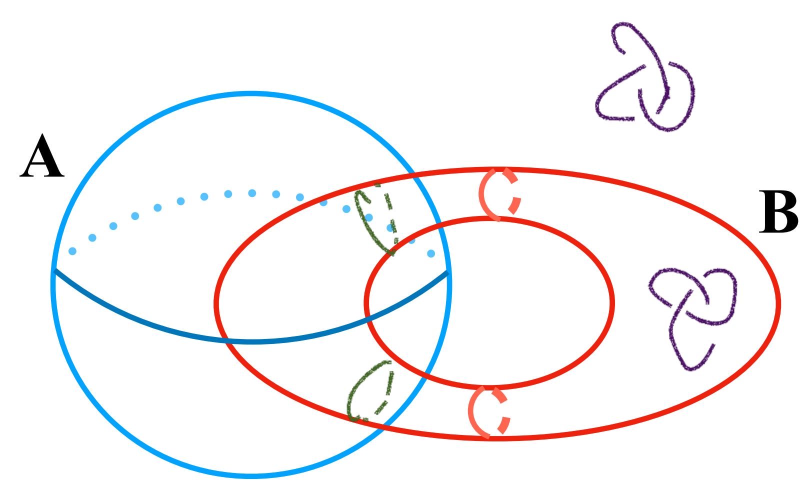

A is not a splitting sphere
Can this happen?

Fundamentalization

Lemma: If L is split then M_L contains a fundamental normal splitting 2-sphere.

Step 2. **Fundamentalize**: If M_L contains a normal splitting 2-sphere then it contains a *fundamental* normal splitting 2-sphere.

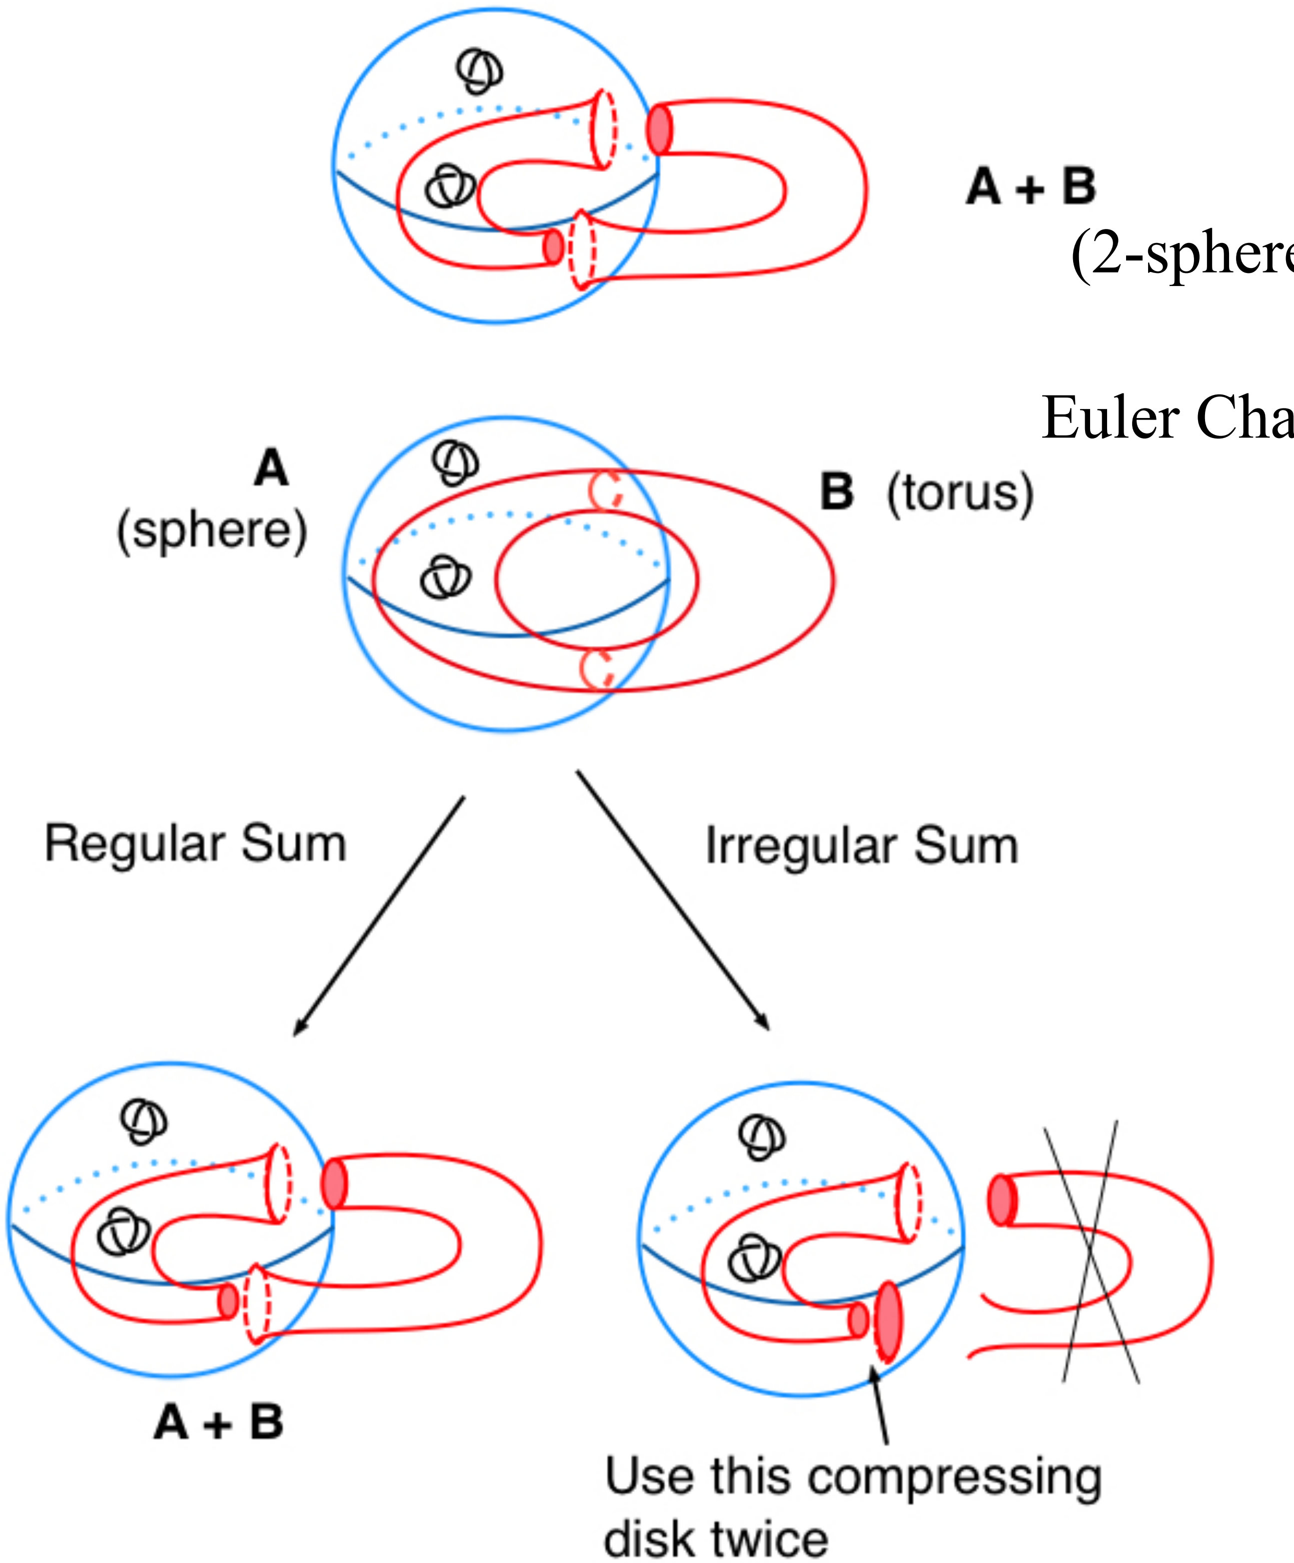

Since the sphere A does not split the link, we can replace one part of A with another without affecting the splitting properties of the resulting surface. In particular, we can use a least weight compressing disk twice.

If **A** is not a splitting sphere then the torus **B** can be surgered to find a splitting 2-sphere of smaller weight than **v.**

But is it normal?

Fundamentalization

Lemma: If L is split then M_L contains a fundamental normal splitting 2-sphere.

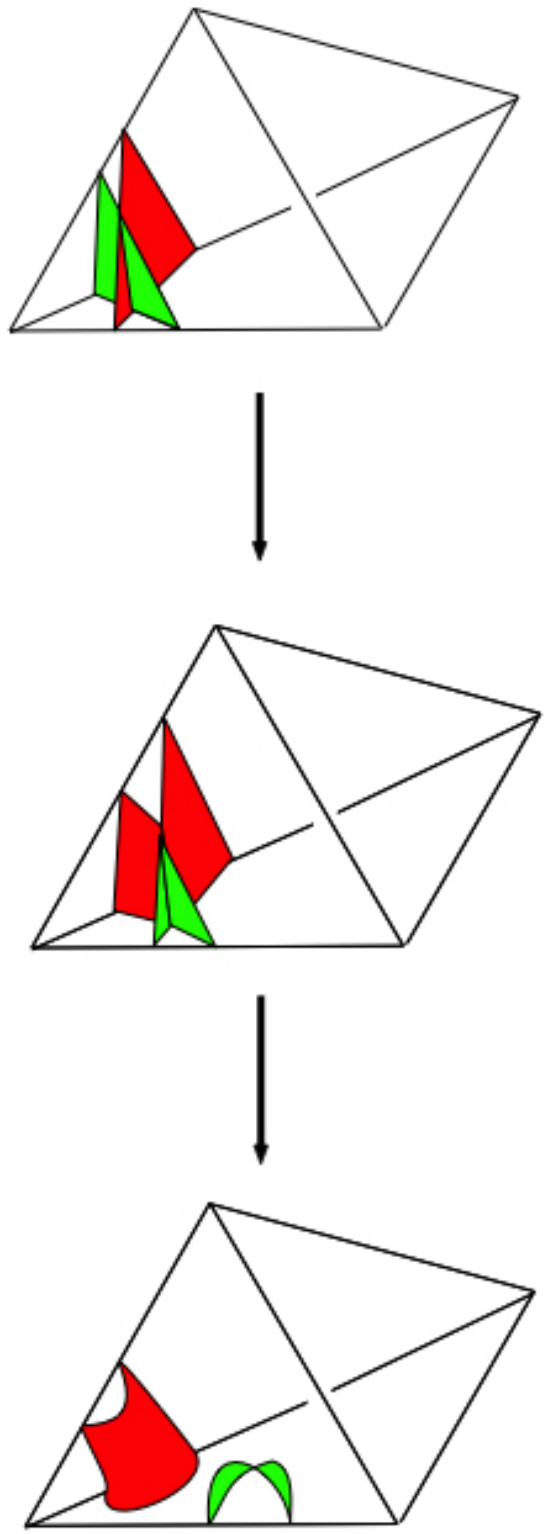

If irregular sum gives a component which is a splitting 2-sphere of smaller weight then **v**, then this component can be normalized to produce a splitting 2-sphere of even smaller weight. This leads to a contradiction.

Irregular sum

If one of these is part of a splitting 2-sphere, it can be normalized to further reduce its weight.

Conclude: A normal splitting 2-sphere of smallest weight is fundamental.

An Algorithm for SPLIT LINK

SPLIT LINK
Instance: A link L in S^3 with complement M_L.
Question: Does M_L contain a 2-sphere that separates the components of L?

The Algorithm

1. Construct all fundamental solutions.
2. Check whether any of the fundamental solutions satisfy $\boldsymbol{\chi}(\mathbf{v}) = 2$ and separates components of L. If one does, then answer "YES". Otherwise answer "NO".

Proof that the algorithm works as claimed:
If there is a splitting 2-sphere then there is a fundamental vector **v** that represents a splitting 2-sphere. A fundamental solution represents a 2-sphere if its Euler characteristic is two. There are finitely many fundamental solutions, and each can be checked to see if it is a 2-sphere that separates components of L.

Corollary: UNKNOTTING