

MATH 697: Topics in Topology

(Winter, 2006)

Set One: Basic Notions

We begin with a discussion of cell-decompositions and triangulations. We assume familiarity with rectilinear triangulations and the basic notions of piecewise linear topology. However, since our presentation incorporates cell-decompositions and more general notions of triangulations, we include definitions, notation and useful facts from this point of view. We have added an Appendix, which provides numerous examples of triangulations and cell-decompositions used to illustrate the theory.

A cell-decomposition of a 3-manifold distinguishes certain surfaces, the normal surfaces relative to that cell-decomposition. They are analogous to stable (minimal) surfaces in geometric analysis and are determined by a stability condition that is equivalent to their meeting the cells of the cell-decomposition in a very special way. However, since the cells in our cell decompositions are not embedded, we need more care in the definition of normal surfaces in this setting. For example, it is traditional to say that a surface in a 3-manifold is normal (relative to a triangulation of the manifold) if and only if it meets each tetrahedron in a collection of normal triangles and normal quadrilaterals. However, for our triangulations, tetrahedra may not be embedded and it is quite possible that a normal surface meets a tetrahedron in an annulus or even a Möbius band; in fact, for a one-tetrahedron triangulation of a 3-manifold, a closed, normal surface might possibly be a sphere, torus or Klein bottle and is contained entirely in the one tetrahedron.

Cell-decompositions and Triangulations.

A cell-decomposition presents a space as a collection of cells with some face identifications. Specifically, suppose $\{\tilde{\Delta}_1, \dots, \tilde{\Delta}_t, \dots\}$ is a pairwise-disjoint collection of compact, convex, linear n -cells. The collection may be infinite. Suppose Φ is a family of affine isomorphisms pairing $(n-1)$ -dimensional faces of the cells in Δ so that if $\phi \in \Phi$, then ϕ is an affine isomorphism from an $(n-1)$ -face $\tilde{\sigma} \in \tilde{\Delta}_i$ to a distinct $(n-1)$ -face $\tilde{\sigma}' \in \tilde{\Delta}_j$, possibly, however, $i = j$. There is an equivalence relation induced on the disjoint union $\Delta = \tilde{\Delta}_1 \amalg \dots \amalg \tilde{\Delta}_t \amalg \dots$ by a family of face-pairings and we use Δ/Φ to denote the space obtained from this equivalence relation with the identification topology. We denote the projection map by $\rho : \Delta \rightarrow \Delta/\Phi$ and incorporate a locally finite condition by requiring that for any point p in Δ/Φ , $\rho^{-1}(p)$ is a finite. We say the family of face-pairings is *monogamous* if a face is paired with at most one other face. If Φ is a family of monogamous face-pairings, then in dimension two ($n = 2$), Δ/Φ is always a 2-manifold. However, for 3-cells, even with the assumption that the family of face-pairings is monogamous, we may not get a manifold (see Figures 2, 3 and 4).

A specific example that fails to be a 3-manifold along an edge is given in Figure 3. Here there is a single face-pairing; the face $(0, 1, 2)$ is paired with the face $(2, 3, 0)$, where the notation indicates that the vertex 0 goes to the vertex 2, the vertex 1

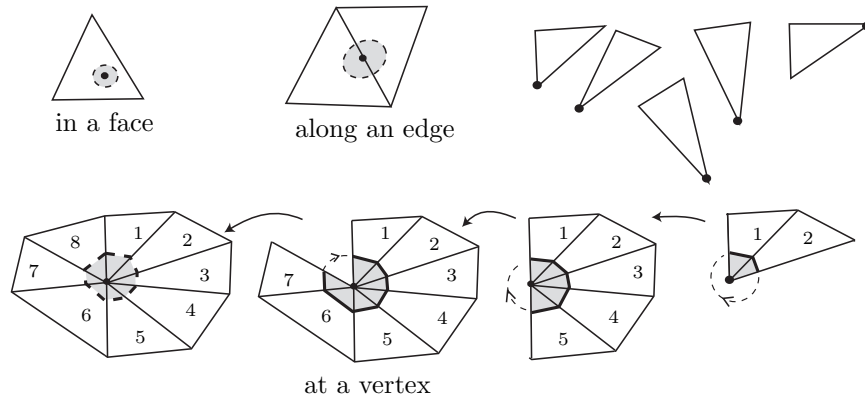


FIGURE 1. Neighborhoods of points for $n = 2$; Δ/Φ is always a manifold.

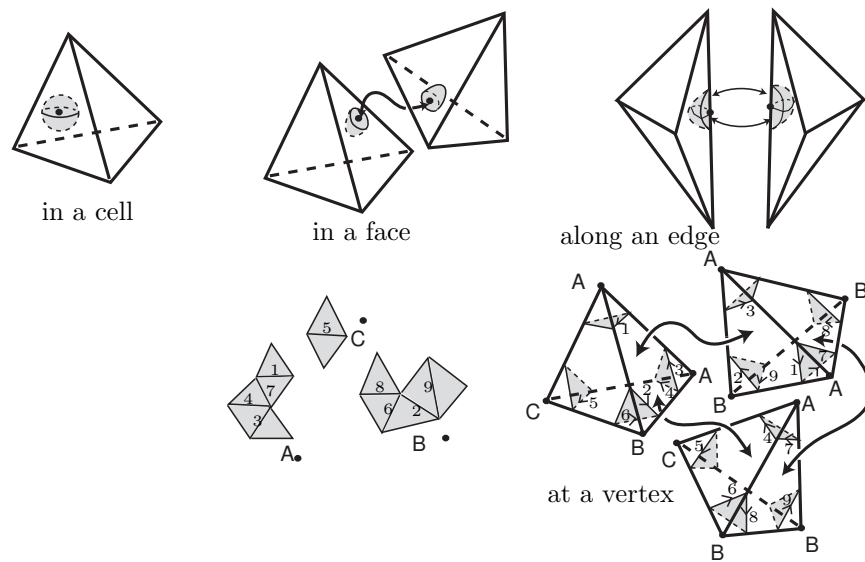


FIGURE 2. Neighborhoods of points for $n = 3$; Δ/Φ may fail to be a manifold along the image of an edge or at the image of a vertex.

goes to the vertex 3 and the vertex 2 goes to the the vertex 0. In this example, after identification, the midpoint of the edge e (the point m in Figure 3) has a small neighborhood in Δ/Φ that is topologically equivalent to the cone over a projective plane. We avoid such identifications, which occur when a small cycle linking an edge is orientation reversing. (Notice if we use oriented cells and orientation-reversing face-pairings, we can avoid the possibility that the identification space, Δ/Φ , might not be a 3-manifold at an image of the midpoint of an edge; however, doing this would exclude non orientable 3-manifolds.) In Figure 4, after identification, the point in Δ/Φ that is the image of the vertices is not a manifold point; a small

neighborhood is the cone over the Klein bottle. We do not, in general, avoid this as such cell-decompositions play a useful role.




FIGURE 3. In general, Δ/Φ may not be a 3-manifold at m , the image of the midpoint of an edge. Here a small neighborhood of the midpoint m , after identification, is the cone over a projective plane. We avoid this phenomenon.




FIGURE 4. In general, Δ/Φ may not be a 3-manifold at the image of a vertex. Here a neighborhood of the image of the vertex is the cone over a Klein bottle.

We collect all this information into a single symbol \mathcal{T} and call \mathcal{T} a *cell-decomposition* of the space Δ/Φ ; in this case, we also use just $|\mathcal{T}|$ to denote the space Δ/Φ . If each cell is a tetrahedron (for $n = 2$ a triangle), we call \mathcal{T} a *triangulation* of the space Δ/Φ .

For a 3-dimensional cell-decomposition \mathcal{T} , a *cell* (*tetrahedron*), *face*, *edge*, or *vertex* in \mathcal{T} is, respectively, the image of a cell (tetrahedron), face, edge, or vertex from the collection $\{\tilde{\Delta}_1, \dots, \tilde{\Delta}_t\}$. We will denote the image of the faces by $\mathcal{T}^{(2)}$, the image of the edges by $\mathcal{T}^{(1)}$ and the image of the vertices by $\mathcal{T}^{(0)}$. We have analogous notions for a 2-dimensional cell-decomposition. We call $\mathcal{T}^{(i)}$ the *i-skeleton* of \mathcal{T} ; but, generally, we just refer to these as the faces, edges or vertices of \mathcal{T} . We will denote the image of $\tilde{\Delta}_i$ by Δ_i and call $\tilde{\Delta}_i$ the *lift* of Δ_i . In the 2-dimensional case, a cell is the quotient of a single cell, an edge is the quotient of one or two edges and a vertex may be the quotient of a number of vertices. For a 3-cell-decomposition, a cell (tetrahedron) is the quotient of a unique cell and a face is the quotient of one or two faces; edges and vertices may be the quotient of a number of edges or vertices, respectively. In the latter case, we define the *order* or *valence* of an edge e of \mathcal{T} to be the number of edges in the collection $\mathbf{\Delta} = \{\tilde{\Delta}_1, \dots, \tilde{\Delta}_t\}$, which are identified

to e . If the link (here we mean the boundary of a small regular neighborhood and not the combinatorial link) of each vertex is either a 2-sphere or a 2-cell, then the underlying point set $|\mathcal{T}|$ is a 3-manifold, possibly with boundary, and we say \mathcal{T} is a *cell-decomposition* of the 3-manifold $M = |\mathcal{T}|$. If each cell in $\mathbf{\Delta}$ is a tetrahedron, then we say \mathcal{T} is a *triangulation* of the 3-manifold $M = |\mathcal{T}|$. As noted above, in the 2-dimensional case, we always get $S = \mathbf{\Delta}/\mathbf{\Phi}$ a 2-manifold. In this situation, if each cell is a triangle, we call the cell-decomposition a *triangulation* of the 2-manifold, S .

In the literature, one is likely to find our notion of a triangulation referred to as a *pseudo-triangulation* and the term triangulation reserved to mean that tetrahedra are embedded and if two simplices meet at all, then they meet in a face of each. We distinguish such triangulations by saying they are *rectilinear triangulations*. A rectilinear triangulation is isomorphic to a subdivision of an Euclidean polyhedron, using standard Euclidean simplicies. Our cells (tetrahedra) are not necessarily embedded; however, the quotient map is an embedding on the interiors of each cell (tetrahedron), face and edge and the intersection of two cells is a union of sub-cells of each. Below we show that the second derived subdivision of one of our triangulations is a rectilinear triangulation; of course, a rectilinear triangulation is one of our triangulations. For a 3-cell-decomposition, if the link of some vertex is a surface, distinct from the 2-sphere or disk, we say \mathcal{T} is an *ideal cell-decomposition* (or *ideal triangulation*) of the 3-manifold $M = |\mathcal{T}| \setminus |\mathcal{T}^{(0)}|$. In this case the vertices of \mathcal{T} are called *ideal vertices*. If the link of an ideal vertex is a closed surface, the *index of the ideal vertex* is the genus of its linking surface. In some cases where we have a mix of genus zero and higher genus ideal vertices, we include the genus zero vertices into our manifold and may not refer to them as ideal vertices. Most often, however, we do not have genus zero vertices in an ideal triangulation. Figure 4 is an example of an ideal triangulation; this is the Gesiking manifold. Other examples are given in the Appendix (see Figures ??, ?? and ??).

In our cell-decompositions, an edge can be a simple closed curve, an edge in a cell with end points (vertices) identified, and faces can take on some interesting configurations. In Figure 5 we give the possibilities for a face which is the image of a triangle. In Figure 5, parts (4) and (5), we have two edges identified to give faces which are *cones* (the latter is a *pinched cone*); in (6), we have a face which is a Möbius band; and in (7) and (8), we have all three edges identified, giving in (7) the classical *dunce hat* (see, for example, Figure 6 (5)) and in (8) a spine for $L(3, 1)$ (see, for example, the classical presentation for $L(3, 1)$ in Figure ?? in the Appendix). In the case of a triangulation of a 2-manifold, we only have those possibilities (1), (2), (3), (4) and (6) in Figure 5 for the triangles. As for tetrahedra, we have in Figure 6 the seven distinct identifications of a single tetrahedron to give an orientable 3-manifold. In Figure 6, the edge e in part (4) has order 1 and in part (5), e has order 5. Any combination of the images of cells, faces, edges or vertices in $\mathbf{\Delta}/\mathbf{\Phi}$ is called a *subcomplex*. We are using closed cells; so, if a cell is included, then all of its subcells are included.

EXAMPLES: Some other examples of interesting of cell-decompositions of 3-manifolds.

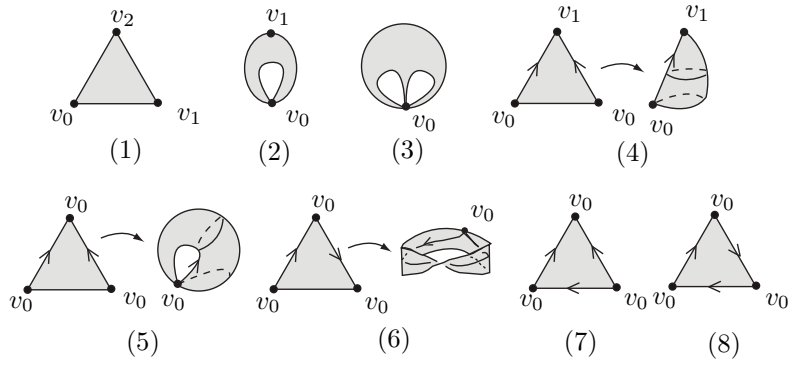


FIGURE 5. Possible configurations for faces in a triangulation.




FIGURE 6. One-tetrahedron triangulations of orientable 3-manifolds.




FIGURE 7. One-cube decomposition of a closed 3-manifold.

0.0.1. Subdivisions of cell-decompositions. Even though we are allowing quite general notions of triangulations and cell-decompositions, we remain in the PL-category and use familiar notions from classical PL-topology, such as subcomplex, subpolyhedron, subdivision, barycentric (derived) subdivision, and regular neighborhood. Several of these notions are considered in the Appendix. Furthermore, we use triangulations as the preferred combinatorial structure for normal surfaces. Even when we use more general cell-decompositions, we restrict ourselves to very special cell-decompositions.

The following is well-known for rectilinear cell-decompositions; here this result requires an interesting method of matching subdivisions (combinatorial structures) in distinct cells along identified faces.

0.0.1. PROPOSITION. *Suppose M is a 3-manifold and \mathcal{C} is a cell-decomposition of M . Then \mathcal{C} can be subdivided to a triangulation, without adding vertices.*

PROOF. Suppose $\{\tilde{\Delta}_1, \dots, \tilde{\Delta}_t\}$ is a pairwise disjoint collection of compact, convex, linear 3-cells. First, we observe we can subdivide each of the cells $\tilde{\Delta}_i$ into tetrahedra without adding vertices. The standard way of subdividing a $\tilde{\Delta}_i$ without adding vertices is to linearly order the vertices of $\tilde{\Delta}_i$; then in each face of $\tilde{\Delta}_i$ there is a “first” vertex in the given linear ordering. The first subdivision comes in the 2-dimensional faces. Thus if σ is a 2-dimensional face and v_σ is the first vertex in the linear ordering in σ , we cone from v_σ over the edges in the boundary of σ not containing v_σ . If σ has n , ($n \geq 3$) edges, then the subdivision of σ has $n - 2$ triangles. Now, to subdivide a cell $\tilde{\Delta}_i$, we have a first vertex in $\tilde{\Delta}_i$ and we cone from it over the subdivide faces of $\tilde{\Delta}_i$ not meeting this first vertex. Since we have always chosen the first vertex in our ordering in each face, the first vertex in $\tilde{\Delta}_i$ is also the first vertex in each 2-dimensional face containing it. Hence, we have a subdivision (actually, a rectilinear triangulation) of each $\tilde{\Delta}_i$ into tetrahedra without adding vertices.

Of course, now we have the problem of face identifications. It is possible that a face $\tilde{\sigma}$ of $\tilde{\Delta}_i$ is identified with a face $\tilde{\sigma}'$ of $\tilde{\Delta}_j$ and the identification is not an isomorphism. However, the image of the subdivision of the face $\tilde{\sigma}$ is a subdivision of $\tilde{\sigma}'$ in which we have not added vertices.

We step aside and make some useful observations about certain subdivisions of compact, convex, planar polygons and about compact convex linear cells.

First, if Q is a rectilinear planar quadrilateral, then we can subdivide Q , without adding vertices, by adding one of the two possible diagonals as an edge. The operation of going from one of these subdivisions of Q to the other is called a *diagonal flip*. Now, suppose F is surface and \mathcal{P} is a subdivision of F . If e is an interior edge of \mathcal{P} , then there are triangles σ and σ' meeting along e . The union $Q = \sigma \cup \sigma'$ is isomorphic to a rectilinear quadrilateral with diagonal e and anti-diagonal the edge e' . If we swap the diagonal e for the diagonal e' , then we have a new subdivision \mathcal{P}' of F , which we say is obtained from \mathcal{P} by a *diagonal flip*. We have the following well known lemma.

0.0.2. LEMMA. *Suppose F is a compact, convex, planar polygon and \mathcal{P} and \mathcal{P}' are each triangulations of F , obtained without adding vertices. Then there is a sequence of triangulations $\mathcal{P} = \mathcal{P}_0, \mathcal{P}_1, \dots, \mathcal{P}_n = \mathcal{P}'$ of F , where \mathcal{P}_{i+1} is obtained from \mathcal{P}_i by a diagonal flip.*

We provide an example in Figure 8; also, see Exercises 0.7 and 0.9.

Now, suppose \mathcal{P} is a triangulation of the surface F . Suppose $\tilde{\tau}$ is a tetrahedron, \tilde{e} is an edge of $\tilde{\tau}$, and $\tilde{\sigma}$ and $\tilde{\sigma}'$ are the two faces of $\tilde{\tau}$ meeting along \tilde{e} . Suppose e is an edge in \mathcal{P} , which is in the interior of F ; hence, there are two triangles σ and σ' in \mathcal{P} meeting along the edge e . If we attach $\tilde{\tau}$ to F taking \tilde{e} to e and extend so that $\tilde{\sigma}$ and $\tilde{\sigma}'$ are attached to the two triangles σ and σ' , respectively, then we get a complex having a copy of the surface F with the triangulation \mathcal{P} along with a copy of the surface F with the triangulation obtained from \mathcal{P} via a diagonal flip in the quadrilateral $\sigma \cup \sigma'$. There are two ways to attach $\tilde{\Delta}$ depending on the direction we attach \tilde{e} to e (we have indicated that we identify $\tilde{\sigma}$ and $\tilde{\sigma}'$ to the two triangles σ and σ' , respectively) but the results are isomorphic. We call this operation a *Pachner move of type $2 \leftrightarrow 2$* on \mathcal{P} . Again, see Figure 8.

We now have the following lemma.

0.0.3. LEMMA. *Suppose $\tilde{\Delta}$ is a compact, convex, linear 3-cell and F is a face of $\tilde{\Delta}$. Furthermore, suppose \mathcal{T} is a triangulation (not necessarily rectilinear) of $\tilde{\Delta}$ obtained without adding vertices and \mathcal{P} is the induced triangulation on F . Let \mathcal{P}' be any triangulation of the compact, convex polygon F obtained without adding vertices. There is a sequence of triangulations $\mathcal{T} = \mathcal{T}_0, \mathcal{T}_1, \dots, \mathcal{T}_n = \mathcal{T}'$ of $\tilde{\Delta}$ so that for \mathcal{P}_j the triangulation induced on F by \mathcal{T}_j , then \mathcal{T}_{j+1} is obtained from \mathcal{T}_j by a Pachner move of type $2 \leftrightarrow 2$ on \mathcal{P}_j and $\mathcal{P}_n = \mathcal{P}'$.*

Now, we are ready to complete the proof of our proposition. If the face $\tilde{\sigma}$ of $\tilde{\Delta}_i$ is identified with a face $\tilde{\sigma}'$ of $\tilde{\Delta}_j$ and the identification is not an isomorphism, then the identification induces a triangulation, say \mathcal{P}' on $\tilde{\sigma}'$ without adding vertices. It follows from Lemma 0.0.3 that we can modify the triangulation on $\tilde{\Delta}_j$ without adding vertices or changing it on any face other than $\tilde{\sigma}'$ so that the face identification is an isomorphism. Hence, it follows we can further subdivide the cells $\{\tilde{\Delta}_1, \dots, \tilde{\Delta}_t\}$ without adding vertices so that we have triangulations and all the face identifications are isomorphisms. \square

In the previous situation, we like to think of having combinatorial structures (triangulations) on each of the cells $\{\tilde{\Delta}_1, \dots, \tilde{\Delta}_t\}$. However, when we go to make face identifications, the combinatorial structures induced on the faces to be identified do not match. The process of changing the combinatorial structure on one of the cells, along a face which is identified, gives a sequence of triangulated surfaces in the cell, each a subdivision of that face and each differing from its predecessor in the sequence of triangulations by a diagonal flip. These are pleated surfaces through which we bring the combinatorial structure of one cell into agreement with the combinatorial structure of the other cell along the face identifications.

0.0.2. Barycentric derived subdivisions. A *barycentric derived subdivision* of \mathcal{T} is formed by taking the barycentric derived subdivision in each of the tetrahedra in the collection $\{\tilde{\Delta}_1, \dots, \tilde{\Delta}_t\}$ and using the induced face identifications. Since face identifications are affine isomorphisms they preserve barycenters and so they, along with the identity maps on simplices of the subdivision interior to simplices in \mathcal{T} , induce face identifications on the collection of tetrahedra resulting from barycentric subdivision of the tetrahedra $\{\tilde{\Delta}_1, \dots, \tilde{\Delta}_t\}$. Following standard notation, we will denote the barycentric derived subdivision of \mathcal{T} by \mathcal{T}' . If $\mathcal{T}'' = (\mathcal{T}')'$ is the barycentric derived subdivision of \mathcal{T}' , we say \mathcal{T}'' is the *second barycentric derived subdivision* of \mathcal{T} , etc. If we denote the barycenter of the simplex $\tilde{\sigma}$ in $\tilde{\Delta}_i$

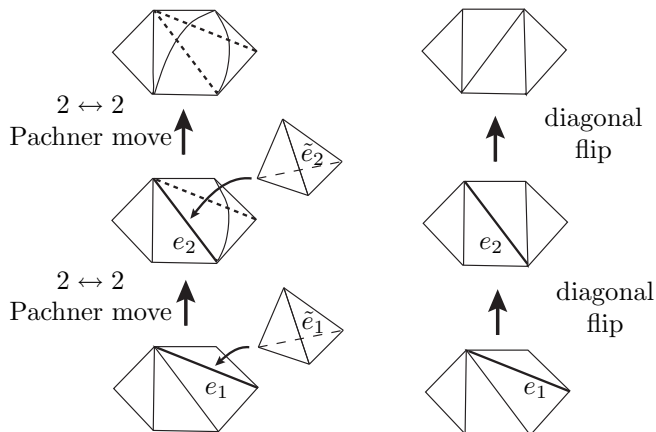


FIGURE 8. Example of two diagonal flips taking one triangulation of the hexagon to another, along with $2 \leftrightarrow 2$ Pachner moves realizing the same diagonal flips.

by $v_{\tilde{\sigma}}$, then the vertices $v_{\tilde{\sigma}_1}, \dots, v_{\tilde{\sigma}_n}$ ($n \leq 4$) form a simplex in \mathcal{T}' if and only if $\tilde{\sigma}_1 < \dots < \tilde{\sigma}_n$, where $\tilde{\sigma}_i < \tilde{\sigma}_j$ means $\tilde{\sigma}_i$ is a proper face of $\tilde{\sigma}_j$.

First, we make some observations in our situation.

Each closed tetrahedron in the barycentric derived subdivision of the collection $\{\tilde{\Delta}_1, \dots, \tilde{\Delta}_t\}$ is embedded under the quotient map. For if this were not the case, there would be two distinct vertices of some tetrahedron in the barycentric subdivision $\tilde{\Delta}'_i$ of $\tilde{\Delta}_i$ identified under the quotient map. But this is impossible, since a tetrahedron in $\tilde{\Delta}'_i$ has the vertices $v_{\tilde{\sigma}_{i_1}}, v_{\tilde{\sigma}_{i_2}}, v_{\tilde{\sigma}_{i_3}}, v_{\tilde{\sigma}_{i_4}}$ where $\tilde{\sigma}_{i_1} < \tilde{\sigma}_{i_2} < \tilde{\sigma}_{i_3} < \tilde{\sigma}_{i_4}$. Hence, in particular, the vertex set of any simplex in \mathcal{T}' is embedded under the quotient map.

If σ' is an n -simplex of \mathcal{T}' , then there is a unique n -simplex σ of \mathcal{T} that contains σ' ; furthermore, σ' contains the vertex v_σ , which is the barycenter of σ . In any lift $\tilde{\sigma}$ of σ , we have a lift $\tilde{\sigma}'$ of σ' and $\tilde{\sigma}'$ is the join from $v_{\tilde{\sigma}}$ to an $n-1$ -simplex in the barycentric subdivision of the boundary of $\tilde{\sigma}$.

The following is a very useful observation and allows us to draw upon familiar notions from classical piecewise linear topology.

0.0.4. PROPOSITION. *Suppose \mathcal{T} is a triangulation of the 3-manifold M . Then the second barycentric derived subdivision of \mathcal{T} is a rectilinear triangulation of M .*

PROOF. Let \mathcal{T}' and \mathcal{T}'' denote the 1st and the 2nd barycentric derived subdivisions of the triangulation \mathcal{T} , respectively.

We want to show that \mathcal{T}'' is isomorphic to a rectilinear triangulation.

First, we give criteria for recognizing when a triangulation is isomorphic to a rectilinear triangulation.

- A triangulation is (isomorphic to) a rectilinear triangulation if and only if
 - the quotient map is a homeomorphism on each closed tetrahedron and
 - if the images of two (closed) simplices meet at all, then they meet in the image of a single closed sub-simplex of each.

or equivalently,

- A triangulation is (isomorphic to) a rectilinear triangulation if and only if whenever a set of vertices in the quotient space determine a simplex, then they determine a unique simplex.

We use the second characterization above to conclude that \mathcal{T}'' is (isomorphic to) a rectilinear triangulation.

Suppose two simplices σ_1'' and σ_2'' in \mathcal{T}'' have the same set of vertices. Since, the quotient map on each simplex in a $\tilde{\Delta}_i$ is an embedding, it follows it is an embedding on vertices of such a simplex; therefore, we have that the two simplices σ_1'' and σ_2'' have the same dimension, say n . From the above, there is a unique n -simplex σ_1' of \mathcal{T}' which contains σ_1'' ; furthermore, σ_1'' contains the vertex $v_{\sigma_1'}$, which is the barycenter of σ_1' . In any lift $\tilde{\sigma}_1'$ of σ_1' , we have a lift $\tilde{\sigma}_1''$ of σ_1'' and $\tilde{\sigma}_1''$ is the join from $v_{\tilde{\sigma}_1'}$ to an $n - 1$ -simplex in the barycentric subdivision of the boundary of $\tilde{\sigma}_1'$. Similarly for σ_2'' . However, since the vertex set of σ_1'' is the same as the vertex set of σ_2'' , it follows that $\sigma_1' = \sigma_2'$ and so we may take $\tilde{\sigma}_1' = \tilde{\sigma}_2'$. Hence, each of $\tilde{\sigma}_1''$ and $\tilde{\sigma}_2''$ are joins from $v_{\tilde{\sigma}_1'}$ to an $n - 1$ -simplex in the barycentric subdivision of the boundary of $\tilde{\sigma}_1'$ (*a priori* different $n - 1$ -simplices in the barycentric subdivision of the boundary of $\tilde{\sigma}_1'$). But the quotient map is a homeomorphism on $\tilde{\sigma}_1'$ and since the vertex sets of $\tilde{\sigma}_1''$ and $\tilde{\sigma}_2''$ have the same image, we have that $\tilde{\sigma}_1'' = \tilde{\sigma}_2''$ and so, $\sigma_1'' = \sigma_2''$. \square

EXERCISES:

EXERCISE 0.1. Confirm the list of 3-manifolds given in Figure 6 exhausts the possible compact 3-manifolds obtained from just one tetrahedron.

EXERCISE 0.2. Confirm the 3-manifolds given in Figure 6 are topologically equivalent to the 3-manifolds as claimed in the figure.

EXERCISE 0.3. Give an exhaustive list of ideal triangulations of 3-manifolds obtained from one tetrahedron.

EXERCISE 0.4. The ideal triangulation given in Figure 4 is called the Gesiking manifold. Construct its orientable double cover and show it is the two tetrahedron ideal triangulation of the Figure Eight knot complement in S^3 , which is given in Figure ?? in the Appendix.

EXERCISE 0.5. Show the minimal number of triangles necessary to triangulate a regular n -gon, in the plane, is $n - 2$.

EXERCISE 0.6. The number of distinct minimal triangulations of a regular $(n + 2)$ -gon, in the plane, is the n^{th} Catalan number. Using this as the definition of the n^{th} Catalan number, calculate the n^{th} Catalan number.

EXERCISE 0.7. Provide a proof of Lemma 0.0.2. First assume that both triangulations are rectilinear, then provide a general proof where one of the triangulations is only isomorphic to a rectilinear one.

EXERCISE 0.8. How many triangles are necessary in a minimal triangulation of a closed 2-manifold? What about a 2-manifold with boundary? Conclude that

a minimal triangulation of a closed 2–manifold has only one-vertex, unless it is S^2 , which has three vertices, or $\mathbb{R}P^2$, which has two. Conclude that a minimal triangulation of a 2–manifold with boundary, other than the disk, must have all its vertices in the boundary and then just one in each boundary component. What are the two possibilities for a minimal triangulation of a disk?

EXERCISE 0.9. Show the two minimal triangulations of the 2–sphere differ by a diagonal flip.

EXERCISE 0.10. (OPEN) How many distinct minimal triangulations are there of a closed, orientable 2–manifold of genus g . This is known for $g = 1, 2, 3$; however, it is not known for $g \geq 4$. In fact, there are no good bounds known for this number if $g \geq 4$.

EXERCISE 0.11. What is the minimal number of tetrahedra needed to triangulate each of the regular solids? How many distinct such triangulations are there for each of the regular solids?

EXERCISE 0.12. (HARD) List the closed, orientable 3–manifolds obtained from face identifications on each of the regular solids. Figure 6 includes such a list for the tetrahedron. As a warm-up, list the closed, orientable 3–manifolds obtained from face identifications of the cube.

EXERCISE 0.13. (HARD) Show that a minimal triangulation of a closed, orientable, irreducible 3–manifold has just one vertex, unless the manifold is one of $S^3, \mathbb{R}P^3$ or $L(3, 1)$. S^3 has two one-tetrahedron triangulations, one with one vertex and one with two vertices; $\mathbb{R}P^3$ has two two-tetrahedron triangulations, one has one vertex and one has two vertices; and $L(3, 1)$ has four two-tetrahedron triangulations, two have one vertex and two have two vertices. Give these examples.

0.1. Normal Surfaces

Given a cell-decomposition \mathcal{C} of the 3–manifold M , we define a normal surface in M (relative to \mathcal{C}) by the way it meets the cells in \mathcal{C} . We provide a number of examples from this point of view in terms of our triangulations. However, normal surfaces can just as well be defined as stable surfaces relative to certain normalization (or shrinking or complexity decreasing) moves relative to the cell-decomposition. This latter view is analogous to stable or minimal (area) surfaces in geometric analysis. The two notions will be shown to be equivalent in Chapter 2, where we also use shrinking to establish the existence of many normal surfaces.

We generally use normal surface theory only in triangulations. Even when we use more general cell-decompositions (truncated tetrahedra, truncated prisms, regions between parallel quads and parallel triangles), we restrict our considerations to very special types of normal surfaces.

0.1.1. Definitions. Suppose \mathcal{C} is a cell-complex and $X = |\mathcal{C}|$ is the underlying topological space. An isotopy of X is a *normal isotopy* (with respect to \mathcal{C}) if it leaves the various cells of \mathcal{C} invariant.

A properly embedded arc in a compact convex linear cell is called a *normal arc* if its end points are in $(n - 1)$ –dimensional faces of the boundary and it is normally isotopic to a straight line segment. In particular, it follows that a normal arc has its end points in distinct $(n - 1)$ –faces. A curve in a compact, convex linear cell is

called a *normal curve* if it is normally isotopic to a curve, which meets each cell in a collection of normal arcs. The elementary components of normal surface theory are determined by special normal curves in the boundaries of the 3-cells of our cell-decomposition. In particular, we call a properly embedded disk in a compact, convex, linear 3-cell an *elementary disk* if it is normally isotopic to a properly embedded disk whose boundary is a normal curve which meets no edge in \mathcal{C} more than once. Up to normal isotopy there are only finitely many equivalence classes of elementary disks in a compact, convex, linear 3-cell; these equivalence classes are called *elementary disk types*.

Suppose \mathcal{T} is a cell-decomposition or ideal cell-decomposition of the 3-manifold M and S is a properly embedded surface transverse to the 2-skeleton of \mathcal{T} . Suppose c is a component of S in the cell Δ_i . Then c is the image of a properly embedded surface, \tilde{c} in $\tilde{\Delta}_i$. We will call \tilde{c} the *lift* of c .

Now, if \mathcal{T} is a cell-decomposition of the 3-manifold M , we say a surface F is a *normal surface* in M (with respect to \mathcal{T}) if F meets each cell of \mathcal{T} in the images of a collection of elementary disks in the cells of $\mathbf{\Delta} = \{\tilde{\Delta}_1, \dots, \tilde{\Delta}_n, \dots\}$. That is, the surface F is a normal surface if and only if the lift of every component of F in a cell of $\mathbf{\Delta}$ is an elementary disk. The elementary disks of normal surface theory for triangulations are the *normal triangles* and *normal quadrilaterals (normal quads)* in a tetrahedron. The normal triangles and normal quads are shown in Figure 9. There are four types of normal triangles and three types of normal quads in each tetrahedron (no identification). In Figure 10, we give some examples of elementary disks in cells which are truncated tetrahedra, truncated-prisms, a cube and an octahedron. In a truncated tetrahedra or truncated prism, besides normal triangles and normal quads, which are also normal in the tetrahedron before truncation, one may have other elementary disks. For example, in Figure 10, we show a normal hexagon and a normal octagon in a truncated tetrahedra and some normal quads in a truncated prism; we also have quadrilaterals, hexagons and octagons in a cube and in an octahedron. This generality exhibited in the elementary disks in cells other than tetrahedra brings complication to the combinatorics and is a major reason we generally use triangulations. In Figure 11, we show some examples of normal surfaces in triangulations of 3-manifolds. There is a 2-sphere made up of two normal triangles (Figure 11(1)), a normal torus, which is a genus one Heegaard torus for S^3 , made up of a single normal quad (Figure 11(2)), and a Möbius band, which is properly embedded in a solid torus and is made up of a single normal quad (Figure 11(3)). In Figure 12 and 13, we give examples of normal surfaces in a layered triangulation of a solid torus. In Figure 12A, we have the meridional disk; in B we have the vertex-linking disk; in Figure 13A we have a Möbius band and in B we have a once-punctured Klein Bottle.

In a tetrahedron in M , after identifications, a normal triangle can take on any of the possible orientable identifications of a triangle shown in Figure 5, Parts (1)–(6). A normal triangle can identify to a Möbius band but then the resulting identification space is not a manifold at the associated vertex. It can be made into an ideal triangulation of a manifold; but the manifold is nonorientable. Similarly, a normal quadrilateral, after identification, can take on numerous forms, including some which are nonorientable. The normal disk types give a normal surface F a cell-decomposition made up of elementary disks (normal quads and normal triangles for a triangulation); we call this the *cell-decomposition induced on F* (or the induced

cell-decomposition). Finally, we note that an embedded normal surface must be properly embedded.




FIGURE 9. Examples of elementary disk types (normal triangles and normal quads for a triangulation).




FIGURE 10. Examples of normal disk types in cells other than tetrahedra.

In a cell-decomposition or an ideal cell-decomposition of the 3-manifold M , the boundary of a small regular neighborhood of a vertex is normally isotopic to a normal surface, each component of which is called a *vertex-linking surface*. If \mathcal{T} is a cell-decomposition of the 3-manifold M , a vertex-linking surface is either a disk (the vertex is in ∂M) or a 2-sphere (the vertex is in $\overset{\circ}{M}$, the interior of M). If \mathcal{T} is an ideal cell-decomposition and v is an ideal vertex, a vertex-linking surface about v is a closed 2-manifold possibly having genus $g \geq 1$, where g is the index of the vertex v . In this case we sometimes refer to the vertex-linking surface as a *surface-at-infinity*. If \mathcal{T} is a triangulation or an ideal triangulation, the entire collection of elementary triangles form an embedded, normal surface, each component of which is a vertex-linking surface. The elementary disk types in a vertex-linking surface of a more general cell-decomposition do not allow such a simple combinatorial description. The vertex-linking surfaces give examples of normal surfaces (see, for example, a vertex-linking Klein bottle in Figure 4 and vertex-linking tori in Figures ?? and ?? given in the Appendix).




FIGURE 11. Examples of normal surfaces in triangulated 3-manifolds.

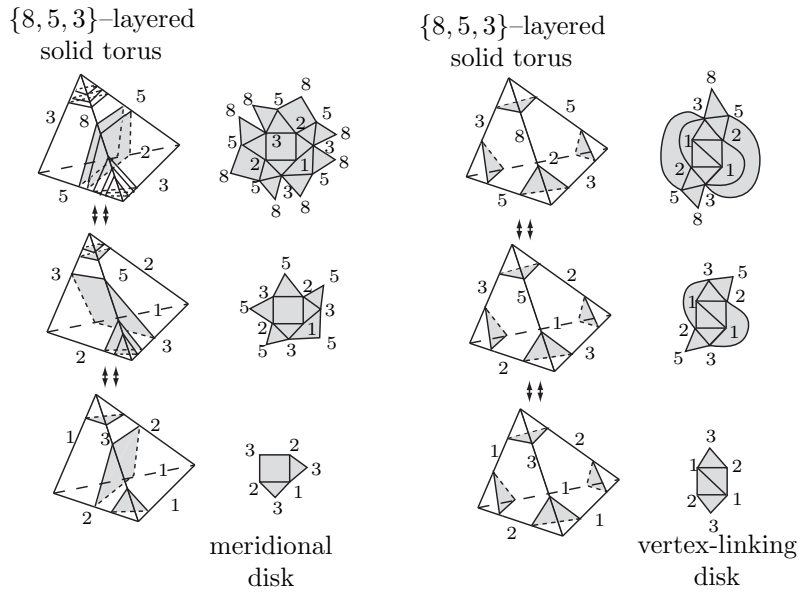


FIGURE 12. Examples of meridional disks and vertex-linking disks in a layered triangulation of a solid torus.

If S is a properly embedded surface in a 3-manifold M and $N(S)$ is a small regular neighborhood of S , the manifold $M' = M \setminus \overset{\circ}{N}(S)$ is said to be obtained from M by *splitting along S* . If S is one-sided in M , then there is a copy, say S' , of a double cover of S (the orientable double cover of S if S is nonorientable and in an orientable 3-manifold) in $\partial M'$. If S is two-sided, then there are two homeomorphic copies, S' and S'' , of S in $\partial M'$. They are in the same component of M' if and only if S does not separate M . If S is a normal surface, then we choose $N(S)$ in such a

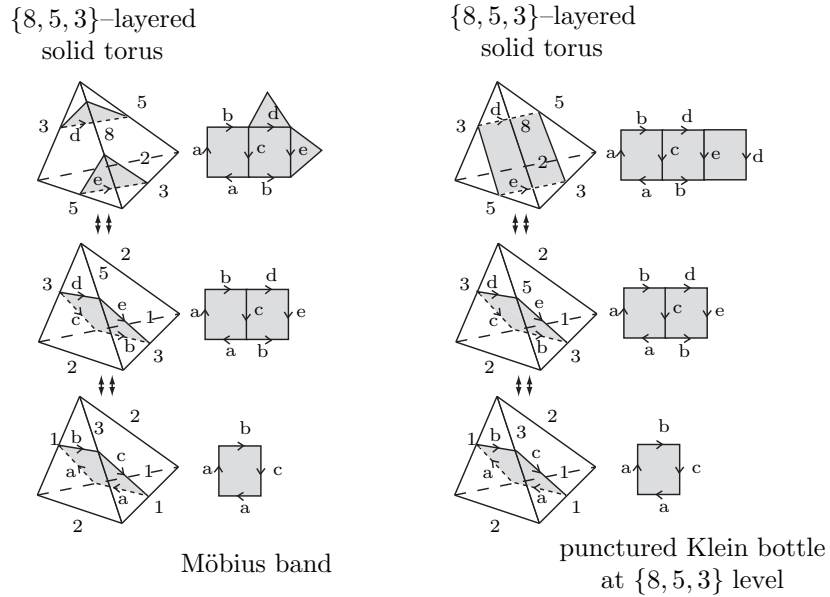


FIGURE 13. Examples of a Möbius band and a punctured Klein bottle in an $\{8, 5, 3\}$ -layered triangulation of a solid torus.

way that the components of its frontier, S' and S'' , are normally isotopic to S (or if S is one-sided, then $S' = 2S$ and has twice the number of elementary cells as S).

If \mathcal{T} is a triangulation of the 3-manifold M , S is a normal surface and M' is obtained by splitting M along S , there is a natural cell-decomposition on M' . A cell in this decomposition is just a component of a $\tilde{\Delta}_i$ split along the normal disks in the lifts of S and the face identifications are just the face identifications for \mathcal{T} restricted to the faces of our new cells. We call this the *induced cell-decomposition* on M' , the manifold obtained by splitting M along S . Let \mathcal{C} denote the induced cell-decomposition on M' . An elementary disk in a cell in \mathcal{C} , which misses $S' \cup S''$ (so, is a quad or triangle), also is an elementary disk in \mathcal{T} ; however, there may be numerous distinct disk types in \mathcal{C} , which miss $S \cup S'$, but all represent the same disk type in \mathcal{T} .

Finally, if S is a two-sided normal surface, M' is the manifold obtained by splitting M along S and S' and S'' the copies of S in $\partial M'$, then there is a natural identification of S' and S'' to recover the 3-manifold M with S' and S'' being identified to S in M . We will refer to this as *re-attaching along S' and S''* . In addition, if S^* is a normal surface in the cell-decomposition induced on M' (M split along S) and possibly now S^* meets $S' \cup S''$, then when we re-attach along S' and S'' , we get a piecewise linear subset, denoted $S \cup S^*$, which is the image of $S' \cup S'' \cup S^*$ in M . We will call this the *piecewise linear normal surface* obtained from S and S^* .

EXERCISES:

EXERCISE 0.14. In each of the examples in Figure 6 exhibit the distinct embedded, connected, normal surfaces. (Later there will be an algorithm to do this. At that time you can check your list against that generated by the algorithm.)

EXERCISE 0.15. List all the normal disk types in a cube. More generally, find the number of normal disk types in each of the regular solids.

EXERCISE 0.16. List all embedded, connected, normal surfaces in the layered triangulations $\{7, 5, 2\}$ and $\{5, 4, 1\}$ of the solid torus. There are only finitely many.

EXERCISE 0.17. Suppose \mathcal{C} is a cell decomposition of a surface F and γ is a normal curve in F (with respect to \mathcal{C}). Set $L(\gamma) = |\gamma \cap \mathcal{C}^{(1)}|$, the number of times γ meets the 1-skeleton of \mathcal{C} . We call $L(\gamma)$ the *length of γ* . Consider normal curves on the boundary of a tetrahedron, a triangulation of the 2-sphere. Except for the four, length 3 vertex-linking curves, show that the length of a normal curve is a multiple of 4.

EXERCISE 0.18. (HARD) Consider the four-tetrahedron layered triangulation of the genus 2 handlebody given in Example ?? of the Appendix. Classify the normal disk in this triangulation. Hint: There are infinitely many.

EXERCISE 0.19. Find an ideal triangulation of the interior of the solid torus and exhibit the vertex-linking torus. (HINT: There is a two-tetrahedron ideal triangulation of the interior of a solid torus.)

EXERCISE 0.20. Show that the three-tetrahedron complex given in Figure 14 is a triangulation of the solid torus. Show all embedded normal disk. Later check your answer against that given by the algorithm. (HINT: There are normal disks which are not vertex-linking.)




FIGURE 14. Example of a three-tetrahedron triangulation of a solid torus with two vertices.

EXERCISE 0.21. OTHER GOOD EXERCISES??