Graphs and Surfaces

Rich Schwartz

March 3, 2016

1 Definition of a Surface

There are many definitions of a surface, and they all turn out to be equivalent, provided you define the notion of equivalence correctly. I'll give a definition which is well-adapted to drawing graphs using polygonal edges.

A *convex polygon* in $\boldsymbol{R}^n$ (that's right, $\boldsymbol{R}^n$) is a subset of $\boldsymbol{R}^n$ having the form $T(P)$, where P is a convex polygon in $\boldsymbol{R}^2$ and $T: \boldsymbol{R}^2 \to \boldsymbol{R}^n$ is a *linear map*. Two convex polygons in $\boldsymbol{R}^n$ *meet cleanly* if their intersection is either a single common vertex or else a single common edge. A *surface* is a finite union $P_1 \cup ... \cup P_m$ of convex polygons in $\boldsymbol{R}^n$ such that

- Every two polygons are either disjoint or meet cleanly.
- For each pair (e, P_i), where e is an edge of P_i, there is a unique pair (e, P_j) where e is an edge of P_j and $j \neq i$.
- If some finite collection of the polygons has a common vertex, then these polygons may be arranged in cyclic order so that each two consecutive polygons meet along an edge that is incident to the vertex.

Let's see what this definition means intuitively. Suppose you start walking around on P_1. Eventually you wander over to some edge e_1. Then you can keep going onto a unique new polygon, say, P_2, which also has e_1 an edge. Then you wander around on P_2 and you eventually come to some edge e_2. You don't fall off; you just move onto another polygon P_3, and so on. The surface has no free edges. What does the surface look like around a vertex? Well, there are finitely many polygons which have this vertex, and the fit together like slices of a pizza. So, at every point, the surface looks "surface-like".

Here is an example of an object that satisfies the first two conditions and not the third. You could take two cubes – I mean their boundaries – and make them touch corner to corner. The problem is that the 6 squares having this common vertex come in two 3-cycles rather than a single 6-cycle.

This definition of a surface is quite general, and it makes it easy to discuss drawings of graphs. Since the surface lives in $\boldsymbol{R}^n$, it makes sense to say what a polygonal path on the surface is. So, a drawing of a graph on the surface is just a finite graph contained in the surface, such that every two vertices are joined by a polygonal path, and no two polygonal paths cross.

2 Equivalence of Surfaces

Let S and S' be surfaces. We say that S' is a *refinement* of S if, as a set, $S' = S$, but each polygon of S' is contained in some polygon of S. Intuitively, we get S' just by breaking some of the polygons of S into smaller pieces. For instance, imagine that S is a tetrahedron, and then you get S' by dividing each of the triangles of S into smaller ones, in the pattern shown in Figure 1.

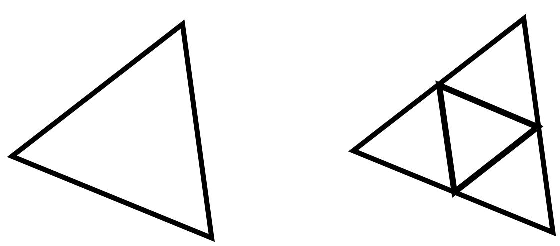

Figure 1: Refining a triangle

An *affine isomorphism* between a polygon $P \subset \boldsymbol{R}^m$ and $Q \subset \boldsymbol{R}^n$ is a linear map $T : P \to Q$ which is a bijection. (The ambient spaces need not have the same dimension.) If P and Q are triangles, then there are always 6 affine isomorphisms between them. In general, there might not be an affine isomorphism at all between P and Q. We say that a *combinatorial bijection* (CB) between surfaces S and T is a bijection $f : S \to T$ such that the restriction of f to each polygon of S is an affine isomorphism onto a polygon of T. The definition is symmetric: The restriction of f^{-1} to each polygon of T is an affine isomorphism onto a polygon of S. Such maps are automatically continuous, but if you are in doubt you can just throw in as an extra hypothesis that both f and f^{-1} are continuous.

Two surfaces S and T are *equivalent* if there are refinements S' and T' of S and T respectively, together with a CB from S' to T'. It is too much to ask that S and T have a CB between them because, in particular, they would have to be composed of the same number of polygons. The definition above is much more flexible. For instance, Figure 2 shows a CB between the refinement of a triangle and the refinement of a square.

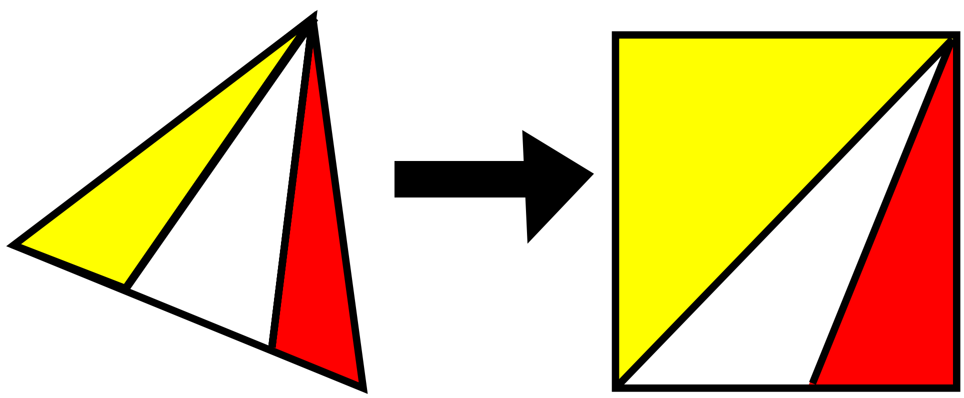

Figure 2: map from a refined triangle to a refined square

The notion of equivalence is pretty flexible. Here is a sample result.

Theorem 2.1 *The boundaries of any two convex polyhedra are equivalent*

Proof: Let A and B be convex polyhedra. Move A and B so that the origin is interior to both. Say that an A *prism* is the set of all vectors λv where $\lambda > 0$ and v is a vector in a face of A. Define a B-prism in the same way. Say that an AB-prism is the intersection of an A prism and a B-prism. Let A' be the refinement of A obtained by intersecting A with all the AB prisms. Likewise define B'. Given an AB prism X, there is an obvious linear isomorphism between $A \cap X$ to $B \cap X$. If you do this map inside each piece, you get the desired map from A' to B'. ♠

Figure 3 illustrates the proof of the above result in the two dimensional case.

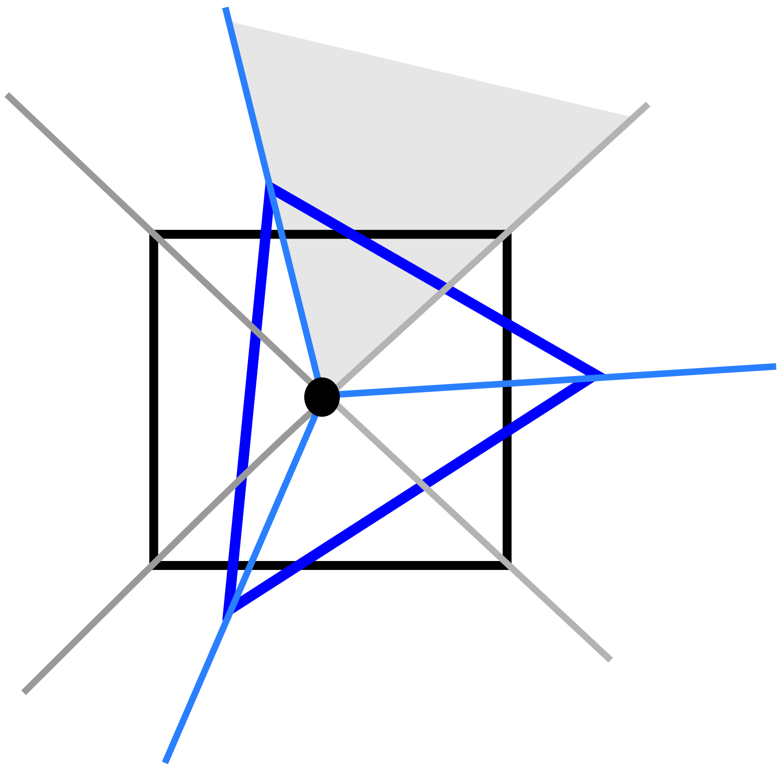

Figure 3: The prism construction

Note that a CB from S' to T' maps polygonal paths to polygonal paths. Hence, the image of a graph drawing on S' is a graph drawing on T'. Also, a graph drawing on S is also a graph drawing on S'. Hence, any graph which can be drawn on one surface can be drawn on any equivalent surface.

3 An Example

Typically we don't worry too much which $\boldsymbol{R}^n$ our surface lives in. We represent surfaces by planar diagrams which indicate how the surface should be assembled in any space which will accept it. Let's do an example. Figure 4 shows a simplified diagram for a surface. We have a square, with the sides identified in the pattern shown.

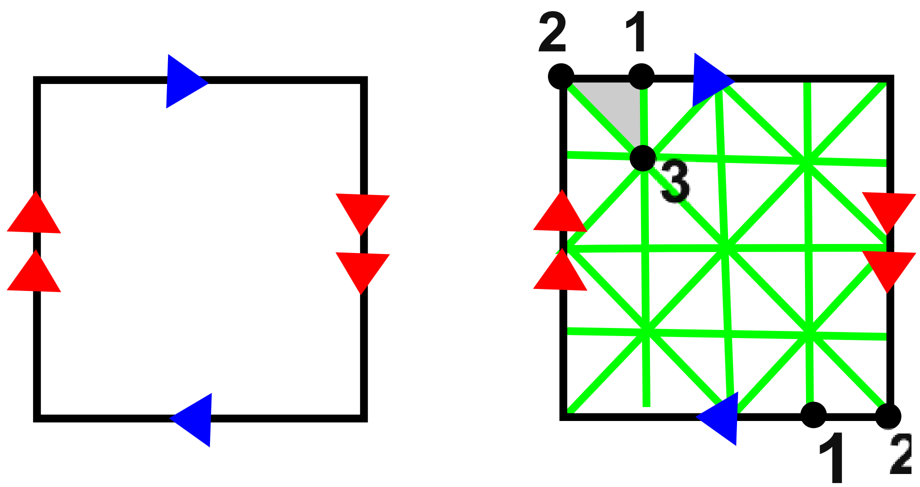

Figure 4: The projective plane and a refinement

The arrows indicate how the sides should be glued together. At first glance, it seems absurd that we could build a surface out of a single square. We aren't allowed to twist the thing around in order to match up the sides. So, to get a surface, we subdivide the square into many small triangles, and then try to embed the triangles into space. The subdivision on the left has 17 vertices. Pick 17 vertices at random in $\boldsymbol{R}^5$ and map each triangle into the convex hull of the corresponding vertices. For instance, the shaded triangle would map to the convex hull of the vertices in $\boldsymbol{R}^5$ labeled 1, 2, and 3. For almost every choice of 17 points in $\boldsymbol{R}^5$, the resulting union of triangles will only intersect at the desired edges. This gives a realization of the projective plane as a surface in $\boldsymbol{R}^5$.

If we picked some different triangulation, we would get a different realization, but the two surfaces would be equivalent. So, the triangulation doesn't really matter much and it is easier to leave it off. When we draw the projective plane or other surfaces we usually leave off the implied subdivision because any subdivision leads to the same surface up to equivalence.

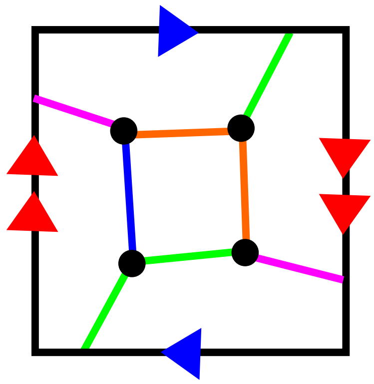

Figure 4: A graph drawn in the projective plane

Figure 5 shows an example of a graph drawn in the projective plane. Some graphs can be drawn in the projective plane and not in the plane. The Peterson graph is an example. (This is a homework problem.)

4 Euler Characteristic

Every surface S comes with a canonical graph drawn in it: The vertices of this canonical graph are just the vertices of the surface, and the edges of the canonical graph are just the edges of the surface. We define the Euler characteristic of the surface to be

$$\chi(S) = V + F - E, \qquad (1)$$

Where V is the number of vertices, and E is the number of edges, and F is the number of polygons.

If S' is a refinement of S then $\chi(S) = \chi(S')$. This follows from the Euler formula in the plane, but let's work it out in some detail. We number the faces of S and in each face P we examine the effect of removing all the edges and vertices of S' in P. Given the planar Euler formula, this removal process does not change the Euler characteristic. Going face to face, we reduce the graph of S' to the graph of S without changing the Euler characteristic. Figure 5 shows that I mean.

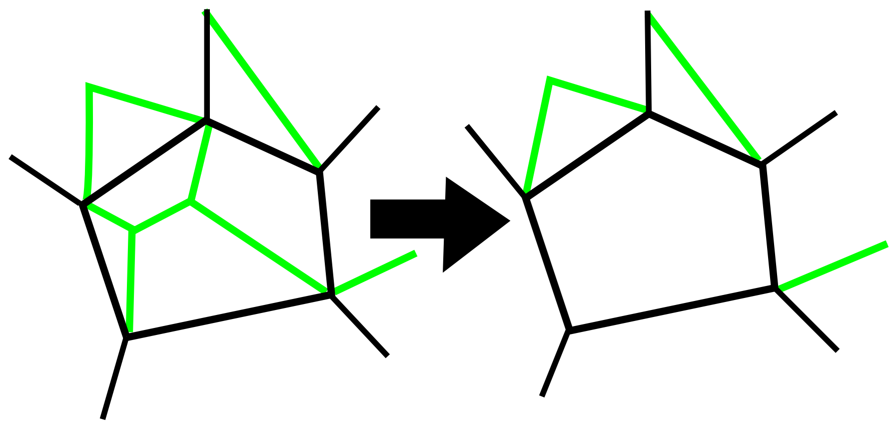

Figure 5: Deleting the extra edges from a face

Note that a BC between two surfaces is a combinatorial bijection. Hence, the graphs of these two surfaces are isomorphic. Combining this with what we already know, we have proved that equivalent surfaces have the same Euler characteristic.

5 The Euler Formula

Let G be a graph drawn in a surface S. A *face* of G is a connected component of $S - G$. We define V, E, F for G just as we did for planar graphs, and we define

$$\chi(G) = V + F - E. \tag{2}$$

The theorem we want to prove is that $\chi(G) = \chi(S)$. However, this need not be true. First of all, as in the planar case, G needs to be connected. Also, we could take the same graph, e.g. a single point, and draw it in any surface. So, we need more structure.

We have only defined the notion of equivalence for surfaces, but we could also define this for faces of G. We say that a face of G is *planar* if that face is equivalent to a triangle. What we mean is that we can subdivide the face into convex polygons, and subdivide the triangle into convex polygons, and then find a CB between then. Figure 6 shows an example. The black edges on the left are part of the graph and the green edges are part of the refinement. Notice that the map does not have to be continuous across the black edges. The map is really just defined on the face and not its closure.

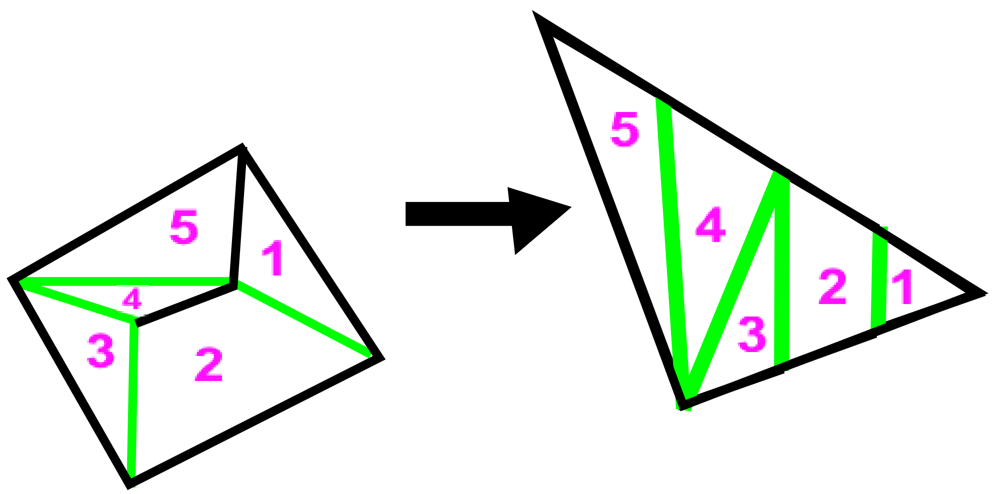

Figure 6: A planar face is equivalent to a triangle

For those of you who know some topology: A face is planar if and only if it is homeomorphic to a disk.

Theorem 5.1 *Suppose that G is a graph in a surface S such that every face of G is planar. Then $\chi(G) = \chi(S)$.*

Proof: Wiggle G a little bit so that G and S have a nice intersection. That is, any two edges intersect in a finite union of points, and no vertex of G is contained in S and no vertex of S is contained in G. Now form a new graph Γ by taking the union of G with all the edges of S and adding intersection points wherever an edge of G intersects an edge of S. You should think of Γ as a kind of common refinement of G and S.

The same argument as above – see Figure 5 – shows that $\chi(\Gamma) = \chi(S)$. We just go around to each face of S and simplify the picture by collapsing edges and then deleting the remaining auxiliary edges. At the same time, the same argument shows that $\chi(\Gamma) = \chi(G)$. For each face of G, we just transplant the picture into a triangle using the equivalence and then run the same argument. ♠

Theorem 5.2 *If G is any graph drawn on S then $\chi(G) \geq \chi(S)$.*

Proof: This result doesn't even require G to be a connected graph. We can add (i.e. draw more) edges until we have a new graph with planar faces. Each time we add edges, we increase E by 1 and may or may not increase F by 1. We end up with a graph G' such that $\chi(G') = \chi(S)$, and $\chi(G) \geq \chi(G')$. ♠

6 Duality

A graph drawn on a surface has a dual graph. The dual graph has one vertex for each face of the original graph, and two vertices are joined by an edge whenever the corresponding faces have an edge in common. Given a graph G, let G^* denote the dual graph.

Theorem 6.1 *If G is a graph drawn on a surface with planar faces, then $G^{**} = G$ and $\chi(G^*) = \chi(G)$.*

Proof: Let V^* denote the number of vertices of G^*. Likewise define E^* and F^*. We have

$$V^* + F^* - E^* \geq \chi(S) = V + F - E.$$

We also know that $E = E^*$ and $F = V^*$. The inequality above tells us that

$$F^* \geq V.$$

On the other hand, the construction gives at least one vertex of G in each face of G^*, because some edges of G cross into each face and they only cross the boundary of the face once. Hence

$$V \geq F^*.$$

Combining these two inequalities, we get $V = F^*$. In other words, G has one vertex for each face of G^*, and two such vertices are joined if and only if the faces of G^* meet across an edge. This shows that $G^{**} = G$. ♠

It also turns out to be true that G^* has planar faces. This follows from the stronger statement that, for a graph Γ drawn in S, we have $\chi(\Gamma) = \chi(S)$ if and only if all faces of Γ are planar. This result is proved by showing that one can always add an edge to a non-planar face without separating it into two faces. But, this requires us to know a bit more about surfaces than we have developed so far. If you are working on a particular surface which you understand well, then you can draw this conclusion by *ad hoc* arguments.

7 A Beautiful Example

Here is a beautiful example of a planar drawing of K_7 on the torus. Consider the usual symmetric tiling of the plane by hexagons. It is possible to color the hexagons in the tilings with 7 colors so that no two neighboring regions have the same color. I'll leave this to you, but I will explain one really great way to see this coloring.

Let $\omega = \exp(2\pi i/3)$ denote the usual cube root of unity. The centers of the hexagons have the form $a + b\omega$ with $a, b \in \boldsymbol{Z}$. At the same time consider the ring $\boldsymbol{Z}/7$. Note that $2^3 = 1$, so in a sense 2 is like a cube root of unity. Define $\phi : \boldsymbol{Z}[\omega] \to \boldsymbol{Z}/7$ by the map $\phi(a + b\omega) = a + 2b$. This gives you a way to assign one of 7 elements to each hexagon center. This is the coloring.

Consider the vectors $V = -2 + \omega$ and $W = 1 + 3\omega$. Note that

$$\phi(V) = \phi(W) = 0.$$

Therefore $\phi(mV + nV) = 0$ for all $m, n \in \boldsymbol{Z}$. This is a way of saying that the coloring forms a regular repeating pattern. We can get a torus by taking the parallelogram whose vertices are

$$0,\ V,\ W,\ V + W$$

and identifying opposite sides. Since addition by V and W preserves the whole tiling, and the coloring, we see that the torus has a tiling by colored hexagons. How many hexagons are in the tiling?

The area of the parallelogram is

$$\mathrm{Im}(V\overline{W}) = 7\sqrt{3}/2.$$

The area of the triangle with vertices $0, 1, \omega$ is $\sqrt{3}/4$. Each hexagons has 6 copies of 1/3 of such a triangle. So, the area of each hexagon is $6 \times 1/3 \times \sqrt{3}/4 = \sqrt{3}/2$. Our torus has 7 hexagons.

We have managed to divide a torus into 7 hexagons, all having a different color. Each hexagon has 6 sides, and touches each of the other hexagons. The dual graph is K_7.

8 Models of Surfaces

Our main result is a little bit unsatisfying because it leave unanswered how to compute the Euler characteristic of a surface. Here we give some nice examples. Let Q be the surface of a cube. We can *attach a handle* to Q by cutting out two little squares of Q and then attaching a cylinder. Call the resulting space Q_1. One should think of Q_1 as a suitcase with a handle. Figure 7 shows a picture of what we have in mind, drawn one dimension down.

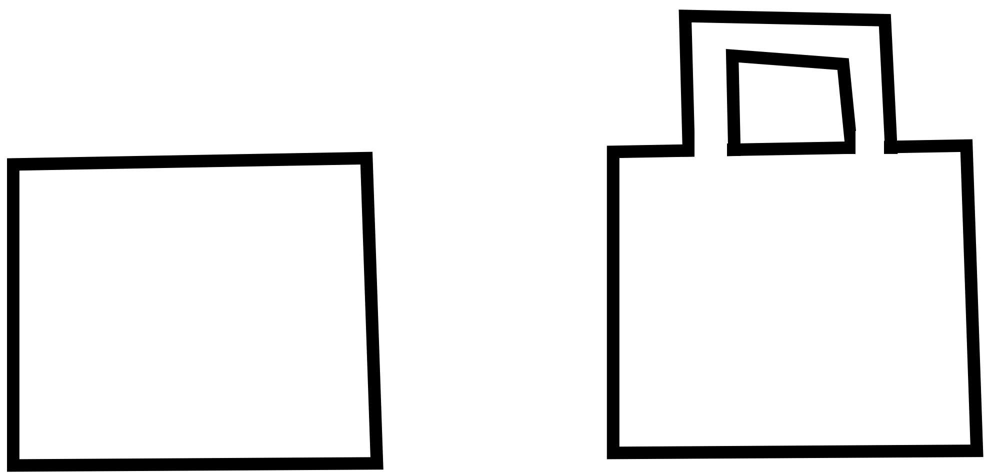

Figure 7: Adding a handle to a cube

We have $\chi(Q) = 8 + 6 - 12 = 2$. We subdivide Q finely so that two of its faces are the pieces we knock out when we attach the handle. Also, the handle itself is equivalent to a sphere with 2 faces knocked out. Therefore

$$\chi(Q_1) = 2 \times \chi(Q) - 4 = 0.$$

We now can form Q_2 by attaching another handle. This time we get

$$\chi(Q_2) = \chi(Q_1) + \chi(Q) - 4 = -2.$$

Every time we attach a handle, we decrease the Euler characteristic by 2. Hence

$$\chi(Q_g) = 2 - 2g.$$

To summarize:

Theorem 8.1 *If G is a graph drawn on Q_g so that all its faces are planar, then $\chi(G) = 2 - 2g$.*

You might wonder: How special are the surfaces Q_g? I'll give an answer without a proof. A surface is called *orientable* if it is possible to orient all the polygons in the surface consistently - i.e. pick a right-hand basis for all of them. Every orientable surface is equivalent to some Q_g.