


Mikhail Lavrov

Start Doing Graph Theory

Part VI: Planar Graphs



available online at <https://vertex.degree/>

Contents

About this document	3
Where it's from	3
What's inside	3
The cover	3
The license	3
21 Planar graphs	4
21.1 Three utilities	4
21.2 Planar graphs	6
21.3 Faces	7
21.4 Euler's formula	10
21.5 Barycentric subdivisions	12
21.6 Technical details	14
21.7 Practice problems	18
22 Planarity testing	21
22.1 Counting edges in planar graphs	21
22.2 Triangulations	23
22.3 Girth and planarity	25
22.4 Subdivisions	27
22.5 Graph minors	29
22.6 Overlap graphs	32
22.7 Practice problems	34
23 Polyhedra	36
23.1 The Platonic solids	36
23.2 Classifying the Platonic solids	38
23.3 Dual graphs	41
23.4 Archimedean solids	43
23.5 Practice problems	46
24 Coloring maps	48
24.1 Coloring maps	48
24.2 The four color theorem	50
24.3 Greedy coloring	52
24.4 Five colors	54
24.5 Hamilton cycles	56
24.6 Coloring empires	57
24.7 Practice problems	60
Bibliography	62

About this document

Where it's from

If you downloaded this PDF file yourself from <https://vertex.degree/contents>, presumably you know what you were doing. But in case you're confused or got the PDF file somewhere else, let me explain!

This is not an entire graph theory textbook. It is one part of the textbook *Start Doing Graph Theory* by Mikhail Lavrov, in case it's convenient for you to download a few smaller files instead of one large file. You can find the entire book at <https://vertex.degree/>; all of it can be downloaded for free.

The complete textbook has some features that the individual parts couldn't possibly have. In this PDF, if you click on a chapter reference from a different part of the book, then you will just be taken to this page, because I can't take you to a page that's not in this PDF file. The hyperlinks in the complete book are fully functional; there is also a preface and an index.

What's inside

In this part of the book, I cover topics related to drawing graphs in the plane. Chapter 21 introduces the fundamental notions, and Chapter 22 presents several ways to determine whether a graph has a plane embedding. Chapter 23 explores the connection between planar graphs and convex polyhedra. Finally, Chapter 24 is about map coloring: the graph coloring problem introduced earlier in the book, applied to planar graphs.

The cover

The cover of this PDF shows a partial coloring of a map which, topologically, can be thought of as a flattened dodecahedron. (The map has 11 regions; the 12th face of the dodecahedron is the outside of the map.)

The license

This document is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License: see <https://creativecommons.org/licenses/by-sa/4.0/> for more information.

21 Planar graphs

The purpose of this chapter

This chapter is an introduction to the ideas you'll need to learn about planar graphs. Most applications of graph theory to geometry are based on Euler's formula and the face length formula, which you will see in this chapter.

The toughest decision for me in presenting this material was to decide how much of the technical geometry to include. The truth is that when working with planar graphs, we don't want to, and also don't have to, worry about these details. The standard academic solution is to begin with a section of technical preliminaries. However, this is also a good way to bore all my readers before we get to the good part.

My compromise was to include rigorous technical proofs of the geometric claims in the last section of this chapter, for the interested reader only. Usually, I cite earlier theorems and lemmas in the textbook when I use them, but in this case, I won't, so that if you don't want to think about these details, you won't have to. However, I have taken care to make sure that every argument in this part of the textbook can be made rigorous using one of the claims in this section.

21.1 Three utilities

Some puzzles are invented because they have a clever solution. Other puzzles are invented because they have no solution; I suspect that sometimes, the motivation is to keep a clever child quietly working on the puzzle for a long time without worries that the child will actually ever solve the puzzle and bother you again.

A classic entry in this genre of puzzles is the three utilities problem [16]. Here, three houses (which we will number 1, 2, and 3) need to be provided water, gas, and electricity by lines from three central utilities (which we will label w, G, and E). The required connections are shown in Figure 21.1a; however, this diagram is not a solution because (for some reason that's never been clear to me) water lines, gas lines, and electricity lines cannot cross. Perhaps the utility providers in this town have not yet learned about the third dimension.

In any case, there is no way to solve the problem. Many attempts look like the diagram in Figure 21.1b; here, curved paths that circle around each other in complicated ways may confuse even the solver into forgetting that house 2 still does not have electricity. (The inhabitants of house 2 are presumably very clear on this point, however.)

But why do all attempts fail? Before presenting you with any general theory, let me give a specific argument for the utility graph. (This is not the most common argument, and [16]

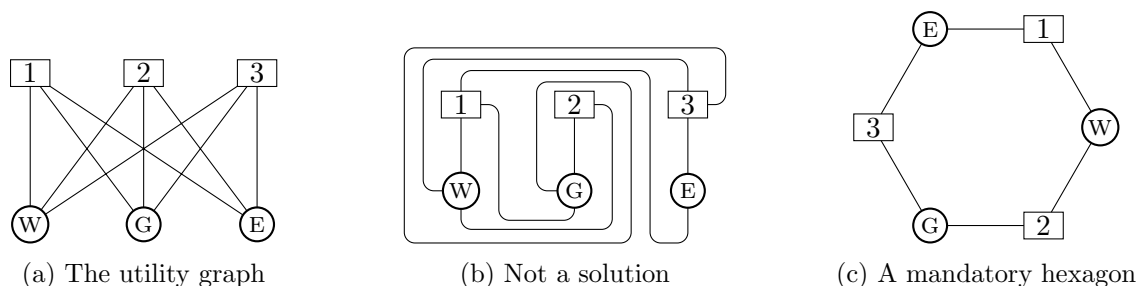


Figure 21.1: The three utilities problem

contains another simple solution. I follow the same approach as West's *Introduction to Graph Theory* [25], and for the same reason: it is a special case of the overlap graph technique we will see in Chapter 22.)

Question: The utility graph shown in Figure 21.1a is isomorphic to another graph we have a name for; what is it?

Answer: We have a couple of names, actually! The utility graph is isomorphic to the complete bipartite graph $K_{3,3}$; it is also isomorphic to the circulant graph $Ci_6(1, 3)$.

The utility graph is Hamiltonian; one possible Hamilton cycle is represented by

$$(1, w, 2, G, 3, E, 1).$$

No matter how this cycle is drawn in a hypothetical solution, it will be a closed loop in the plane with an inside and an outside. (This is even true in Figure 21.1b, though the closed loop is rather complicated!) To help us visualize the closed loop, I will draw it as a regular hexagon in Figure 21.1c, though of course the loop doesn't have to look anything like this.

There are still three edges that we have not drawn: the edges 1G, 2E, and 3W. Only one of these edges can be drawn inside the loop. For example, if we draw edge 1G inside the loop, then it separates vertices 2 and w from vertices 3 and E, so neither 2E nor 3W can also be drawn inside the loop.

Of course, we can still draw either of those edges outside the loop. However, an identical problem occurs there. Suppose we decide to draw edge 2E outside the loop. No matter how we do it, the edge and the loop divide the outside into two pieces: one (call this F_1) bounded by the edges $\{2E, 2G, 3E, 3G\}$ and one (call this F_2) bounded by the edges $\{1E, 1W, 2E, 2W\}$. One of these pieces will be finite and the other infinite, but that doesn't concern us in any way.

What does concern us is that once we've drawn edge 2E outside the loop, we can't do the same thing with edge 1G or 3W. A curve in the plane from 1 to G drawn outside the loop, for example, would start in F_1 and end in F_2 , so at some point it would have to cross one of the boundaries.

Of course, neither drawing 1G inside the loop nor drawing 2E outside the loop is required, but we're stuck no matter what we try: the loop has two sides (inside and outside) and each side only has room for one edge, but we have three edges left to draw. Because $3 > 2$, the three utilities problem has no solution.



Figure 21.2: Planar and non-planar graphs

21.2 Planar graphs

The three utilities problem is one special case of a general problem in graph theory: which graphs can be drawn in the plane without crossings?

Sometimes answering this question is easy. For example, Figure 21.2a shows two ways to draw the complete graph K_4 . In the standard one, the two diagonal edges cross; however, we can draw one of the two edges differently and eliminate that crossing.

Sometimes answering this question is hard. Both of the graphs in Figure 21.2b are drawn with many crossing edges. For one of them, this can be fixed; for the other, there is no way to avoid crossings. However, it is far from obvious which graph has which property; we will need to learn more before we answer this question.

Question: In Figure 21.2a, we avoid crossings by curving one of the edges. However, we can also draw K_4 in the plane with 6 straight edges that don't cross; how?

Answer: Put 3 vertices of K_4 at the corners of a triangle, and put the 4th vertex inside the triangle.

It's important to distinguish between two similar concepts: “a graph that can be drawn in the plane with no crossings” and “a drawing of a graph in the plane with no crossings.” The first of these is a graph invariant. If G and H are two isomorphic graphs, and we can draw G in the plane with no crossings, then we can draw H in the plane with no crossings: just relabel the drawing of G .

However, whether or not a graph is drawn with no crossings is not a graph invariant! Figure 21.2a makes this clear: the graph K_4 can be drawn with crossings, but it can also be drawn without crossings, without changing the graph itself.

Accordingly, we have two definitions to separate these concepts.

Definition 21.1. A *plane embedding* of a graph G is a drawing of G in the plane with no crossings.

A *planar graph* is a graph that has a plane embedding.

This means that our discussion of the the three utilities problem can be taken as a proof of the following proposition:




Figure 21.3: Examples of faces in plane embeddings

Proposition 21.1. *The complete bipartite graph $K_{3,3}$ is not planar.*

Let me emphasize again that a planar graph and a plane embedding are two different things. A planar graph is still just an abstract object: a set of vertices and a set of edges. We haven't picked a particular drawing for it, and there could be many drawings that are different from each other in important ways.

When using the assumption that a graph G is planar in a proof, we should usually begin by choosing a plane embedding of G . However, we should be careful when talking about properties of that plane embedding: they are not properties of the graph G itself, unless we can prove that all plane embeddings of G share those properties. In this chapter, we will see some examples where we can prove this, and some examples where we can't.

We will occasionally have reason to talk about plane embeddings of multigraphs, rather than graphs. This will not come up when deciding if a graph is planar or not: a multigraph is planar if and only if its simplification is planar. But we sometimes want to consider plane embeddings of multigraphs as objects of study in their own right.

21.3 Faces

The diagrams in Figure 21.3 show two plane embeddings in which I've highlighted several regions of the plane. (The shading in regions that I've labeled F_0 in both diagrams is meant to indicate that they are infinite regions, extending outward to the rest of the plane.) We call these the faces of the plane embedding.

Definition 21.2. *The **faces** of a plane embedding are the connected regions of the plane separated by the edges of the plane embedding. One of the faces is infinite, and contains all points sufficiently far from the plane embedding; it is called the **outer face**.*

Faces are the link between the geometry of a plane embedding and the abstract structure of the graph G . In the graph G itself, a region of points does not exist. So what does? What's left is the boundary of the face: the vertices and edges that separate it from the other faces.




Figure 21.4: Boundary walks of the faces in Figure 21.3

In ideal circumstances, the boundary of a face is a cycle, in which case (as usual) it is represented by a closed walk. Don't forget that the cycle doesn't have a prescribed starting point or direction, so it can be represented by many different closed walks.

Question: In Figure 21.3a, which closed walks represent the boundary of the faces?

Answer: The boundary of F_1 is represented by $(1, 2, 4, 1)$, the boundary of F_2 is represented by $(2, 3, 5, 4, 2)$, and the boundary of F_0 is represented by $(1, 4, 5, 3, 2, 1)$.

However, the boundary of a face is not always as nice. In Figure 21.3b, the boundaries of faces F_1 and F_2 are a bit more unusual:

- In the case of F_1 , we would like edge 14 to be taken into account, even though it only separates face F_1 from itself, so the boundary is not a cycle. However, the closed walk $(1, 2, 9, 8, 1, 4, 1)$ still seems to represent the boundary.
- In the case of F_2 , the boundary is not even connected! We need two closed walks: $(2, 3, 0, 9, 2)$ and $(5, 6, 7, 5)$.

In general, the boundary of a face is represented by one or more **boundary walks**. A boundary walk is found by starting at any vertex on the boundary of the face, and following edges out of it in some direction, while staying inside the face. Figure 21.4 shows the boundary walks of the faces of the two example graphs we looked at.

Question: Looking at Figure 21.3b, why does face F_2 have two boundary walks?

Answer: The graph has multiple connected components, each of which contributes a boundary walk to F_2 .

Question: Why is edge 14 in Figure 21.3b on the boundary walk of F_1 twice?

Answer: The same face, F_1 , is on both sides of the edge.

Every edge xy is either on two boundary walks of two different faces, or on the boundary walk of the same face twice. The difference between these two cases comes down to the following test:

Lemma 21.2. *If edge xy is a bridge of planar graph G , then in every plane embedding of G , it appears on the boundary walk of the same face twice.*

If edge xy is not a bridge, then in every plane embedding of G , it separates two faces, and appears once on the boundary walk of each face.

Proof. If the same face F is on both sides of edge xy when it is drawn in the plane embedding, then we can draw a curve from one side of xy to the other while always staying inside F . If xy is deleted, we can complete that curve to a closed loop drawn entirely inside F . Of the two vertices x and y , one is drawn inside that closed loop and the other outside it, and no edges cross the loop; therefore there is no way to get from x to y . This makes xy a bridge.

If edge xy separates two faces F_1 and F_2 , then each of them has a boundary walk using edge xy once. Without loss of generality, such a boundary walk goes from x to y , and returns along some $y - x$ walk without using edge xy again: that walk exists in $G - xy$. Every $y - x$ walk contains a $y - x$ path (by Theorem 3.1), so there is an $x - y$ path in $G - xy$. This means that in G , there is a cycle containing xy , so xy is not a bridge.

The above arguments tell us whether xy is a bridge or not based on what it does in one plane embedding. But regardless of the plane embedding, edge xy either is a bridge in graph G or it isn't, so it must have the same behavior in every plane embedding. \square

The most important property of the boundary of a face is its length. The **length** of a face F , written $\text{len}(F)$, is the sum of the lengths of boundary walks of F . Equivalently, $\text{len}(F)$ is the number of edges on the boundary of F , counting an edge twice if it is not on the boundary of any other face.

When we give the definition of $\text{len}(F)$ in its second form, it seems unnatural and artificial, though you may remember that when we defined the degree of a vertex of a multigraph in Chapter 7, we did something similar. Just like back then, it is all in service of making sure a nice lemma holds in all possible cases:

Lemma 21.3 (Face length formula). *If G is a planar graph with m edges, and a plane embedding of G has faces F_0, F_1, \dots, F_{r-1} , then*

$$\sum_{i=0}^{r-1} \text{len}(F_i) = 2m.$$

Proof. Each edge of G appears once on boundary walks of two faces, or twice on the boundary walks of one face. In either case, it contributes $+2$ to the sum of face lengths, so the contributions of all m edges add up to $2m$. \square




Figure 21.5: Three embeddings of the same graph

The face length formula tells us that the sum of the lengths of the faces does not depend on the plane embedding chosen. (After all, it is equal to $2m$, which certainly cannot change based on the plane embedding!) The individual lengths, however, certainly do depend on the embedding!

Figure 21.5 shows three different plane embeddings of the same 7-vertex graph. (I have deliberately not used consistent colors to discourage you from the temptation of feeling that the faces of one embedding correspond in any way to the faces of another.) However, the face lengths vary:

- The embedding in Figure 21.5a has two faces of length 3, one face of length 5, and an outer face of length 7.
- The embedding in Figure 21.5b has two faces of length 3, and two faces (including the outer face) of length 6.
- The embedding in Figure 21.5c is almost like the one in Figure 21.5a: it also has faces of lengths 3, 3, 5, 7. However, in this case, the embedding of length 5 is the outer face.

21.4 Euler's formula

The face length formula is one of two important identities regarding the faces of a plane embedding. The other one is Euler's formula (one of several mathematical statements by that name!), which helps us count the faces in a plane embedding. You may have noticed that no matter how hard we tried to find different plane embeddings in Figure 21.5, all of them had four faces. This is not a coincidence, and it will follow from Euler's formula that the number of faces in a plane embedding depends only on the graph.

Theorem 21.4 bears Euler's name because he was the first to look at the problem in a general case [6]. Euler did not consider the problem for planar graphs, but for polyhedra, which we will study in Chapter 23. Many other mathematicians later approached the formula from different directions; see [5] for a brief history, as well as further references and many different proofs.

Euler's formula is often stated as $V - E + F = 2$, where V is the number of vertices, E is the number of edges, and F is the number of faces. I can't bring myself to do this, since I think of V and E as sets, not numbers. In graph theory, n and m are the standard variables to use when you want to count the vertices and edges in a graph, respectively. There is no standard

variable for faces, so I will use r (for “region”). Meanwhile, k will be our variable of choice for connected components, as it already was in Chapter 10.

Theorem 21.4 (Euler’s formula). *If a connected plane embedding (of a graph or a multigraph) has n vertices, m edges, and r faces, then*

$$n - m + r = 2.$$

In general, if there are k connected components, this formula becomes

$$n - m + r - k = 1.$$

Proof. We induct on the number of edges, m . Since removing an edge from a graph may disconnect it, we should work directly with the second, more general form of Euler’s formula; the first form will follow by setting $k = 1$.

Our base case is $m = 0$. Here, n (the number of vertices) is equal to k (the number of connected components), and there is only one face: $r = 1$. So $n - m + r - k$ starts out at 1.

Now assume for some $m \geq 1$ that Euler’s formula holds for all plane embeddings with $m - 1$ edges, and consider a plane embedding of a graph G with m edges. We will arbitrarily pick an edge xy to delete.

Question: How can the deletion of edge xy affect n , m , r , and k ?

Answer: It never affects n , and always decreases m by 1.

It may decrease r by 1, if xy previously separated two faces.

It may increase k by 1, if xy was the only way to get from x to y .

We would like $n - m + r - k$ to stay the same, so we would like exactly one one of two things to happen: either r decreases by 1, or k increases by 1, but not both.

Fortunately, this is exactly what Lemma 21.2 tells us! If edge xy is a bridge, its deletion increases k by 1; however, in every plane embedding, the same face is on both sides of xy , so r stays the same. If edge xy is not a bridge, then it separates two faces, so deleting it decreases r by 1; however, it is not a bridge so its deletion does not affect k .

In summary: the number of vertices, edges, faces, and components in the plane embedding of $G - xy$, which we’ll denote (n', m', r', k') , is either $(n, m - 1, r - 1, k)$ or $(n, m - 1, r, k + 1)$. By the induction hypothesis, $n' - m' + r' - k' = 1$, which tells us either that $n - (m - 1) + (r - 1) - k = 1$ or that $n - (m - 1) + r - (k + 1) = 1$; both of these simplify to $n - m + r - k = 1$.

By induction, the formula holds for all plane embeddings. You can see the induction in action in Figure 21.6. Here, we’ve chosen to delete only edges that keep the graph connected, so going from Figure 21.6a to Figure 21.6d, r decreases at each step while k remains equal to 1. At this point, if we delete more edges, it will decrease m and increase k while keeping n and r the same. \square




Figure 21.6: How the triple (n, m, r) changes as we erase edges in a plane embedding

In Euler’s formula, the variables n , m , and k are properties of the graph G , while r is computed from a specific plane embedding. However, we can solve for r in terms of n , m , and k : it is $m - n + k + 1$, or just $m - n + 2$ if G is connected.

Question: How can these observations be reconciled?

Answer: It means that although different plane embeddings can have faces with different shapes and with different lengths, two plane embeddings of the same graph are guaranteed to have the same number of faces: the number predicted by Euler’s formula.

We should still avoid talking about “the faces of G ”, because that is not a well-defined concept, as Figure 21.5 shows. Talking about the number of faces that G has is also incorrect, but if you said it to my face, I would only disapprove a little; after all, if we pick different embeddings, we will still not disagree about the number of faces.

21.5 Barycentric subdivisions

Euler’s formula and the face length formula have applications to geometry in scenarios where the problem can be modeled as a graph. To illustrate this, let me tell you about a problem that once misled me until I thought of applying Euler’s formula to solve it.

The *barycentric subdivision* of a triangle subdivides it into six triangles, by drawing the three medians: the lines connecting each corner of the triangle to the midpoint of the opposite sides. (The three lines in this way always intersect at a single point, which is called the centroid of the triangle.) An example is shown in Figure 21.7b, though the triangle we start with does not have to be equilateral.

We can iterate this process: start with the barycentric subdivision of a triangle, and then take the barycentric subdivision of each of the six small triangles that result. This is shown in Figure 21.7c, but we don’t have to stop there; we can keep going, by taking the barycentric subdivision of each of the 36 small triangles in that diagram. Let’s say that stage n of this process is the result we get when we apply the operation “take the barycentric subdivision of every triangle which does not have any smaller triangles inside it” n times.




Figure 21.7: Iterated barycentric subdivisions

I was interested in the number of vertices in this diagram: I placed a vertex at every point where two or more segments met, or in other words, at every point which is the corner of one or more triangles. If you count the vertices in Figure 21.7, you get 3 vertices in Figure 21.7a, 7 vertices in Figure 21.7b, and 25 vertices in Figure 21.7c. I went one step further and laboriously computed the number of vertices in Stage 3, which turned out to be 121.

I did not even want to think about drawing Stage 4, so my next step was to go to the Online Encyclopedia of Integer Sequences and search for the sequence 3, 7, 25, 121. At the time, there was essentially one candidate formula in the search results (plus a few minor variants which were clearly wrong for this problem), but it was a very nice formula: it suggested that the n^{th} stage would contain $(n + 2)! + 1$ vertices.

Question: At this point, there is a quick way to see that this formula can't possibly be right for all n , because the number of vertices grows at most exponentially. Why?

Answer: The number of triangles in the subdivision is much easier to count: it is 6^n in the n^{th} stage, because each triangle gets subdivided into six triangles. Each triangle has 3 corners, giving $3 \cdot 6^n$ vertices total. This counts some vertices many times, because some vertices are the corners of many triangles; however, it counts each vertex at least once, so there are at most $3 \cdot 6^n$ vertices.

Unfortunately, I only thought of this quick way later. At the time, I spent a while trying to prove the formula $(n + 1)! + 1$ combinatorially. Maybe we should label the vertices with permutations of $n + 2$ elements, but leave the central vertex unlabeled? In the end, it occurred to me to try Euler's formula, and that's when my house of cards fell apart.

Why Euler's formula? Well, the diagrams shown in Figure 21.7 and the ones that come after are all plane embeddings: the vertices are exactly as I have defined them already, and the edges are the line segments connecting the vertices. We can hope that the edges and faces will be easier to count than the vertices, replacing one hard counting problem with two easier ones.

In this case, to apply Euler's formula, we should start with the faces. It is not quite true that there are 6^n faces, all with 3 sides, because there is also an outer face; there are $6^n + 1$ faces total.

Question: What is the length of the outer face?

Answer: It is $3 \cdot 2^n$: in the n^{th} stage, each side of the triangle has been divided into 2^n segments, which are all edges of the plane embedding.

Question: How can we find the number of edges?

Answer: We can add up the lengths of the faces and apply the face length formula.

With 6^n faces of length 3 and one face of length $3 \cdot 2^n$, the sum of face lengths is $3 \cdot 6^n + 3 \cdot 2^n$, so there $\frac{3 \cdot 6^n + 3 \cdot 2^n}{2}$ edges. Solving for the number of vertices in Euler's formula, we get

$$\frac{3 \cdot 6^n + 3 \cdot 2^n}{2} - (6^n + 1) + 2,$$

which simplifies to $\frac{6^n + 3 \cdot 2^n}{2} + 1$. When $n = 4$, the first case I did not consider, this formula tells us that there are 673 vertices. This is the first time the true answer disagrees with the incorrect formula $(n + 2)! + 1$ (which gives 721).

This story has a happy ending, of sorts. Future mathematicians who stumble upon this problem will not have the experience I did, because the correct answer now also shows up in the OEIS: as sequence A380996 [18].

21.6 Technical details

This section contains all the geometric arguments that we will need, in this part of the textbook, to confirm our intuition about how planar graphs can, and cannot be drawn. Though I will not cite specific claims, I have made sure that every geometric fact we need is contained somewhere in this section.

First of all, we should be more precise about what a plane embedding exactly is. There is not much to say about vertices: every vertex is drawn as a point in the plane. An edge should be drawn as some path between its two endpoints; two edges do not contain any common interior points, and in particular their interiors do not pass through any vertices. But what kind of paths should we allow?

In principle, any kind of continuous simple curve will do: the image of a continuous injective function $f: [0, 1] \rightarrow \mathbb{R}^2$, where the points $f(0)$ and $f(1)$ are the two endpoints of the edge. However, studying such things will drown us in a sea of topology for no practical benefit.

It is sufficient to consider edges which are *polygonal curves*: made up of finitely many line segments joined end to end (and not intersecting otherwise). In a closed polygonal curve, or *polygon*, the last line segment ends where the first segment starts. Polygonal curves and polygons can approximate any crazier shape arbitrarily well; you won't be able to tell the difference. Practically speaking, when you encounter a planar graph (for example, derived from a map specified by latitude and longitude coordinates), it will already have edges in this form. Paths and cycles in a planar graph also become polygonal curves and polygons, respectively, in a plane embedding.

To define faces in a plane embedding, we want to be able to tell which points are in the same connected region of the plane (separated by the edges). Our test for this also uses polygonal curves: two points x and y are in the same face if there is a polygonal curve from x to y that does not intersect the plane embedding.

Question: Why is this notion of connectedness practically useful to us?

Answer: It tells us which new edges we could add to the plane embedding, because the edges are also polygonal curves.

There are a few more things to say about this notion of connectedness. From the point of view of theory, the most fundamental is the Jordan curve theorem, which states an obvious-seeming fact: every polygon separates the plane into exactly two connected regions, an inside and an outside, with the polygon as their boundary. (That is, points on the polygon are exactly the points arbitrarily close to both the inside and the outside.) The Jordan curve theorem holds for all continuous curves, not just polygons, but even in the polygonal case it is not easy to prove; I will skip it, to avoid too lengthy a geometric detour.

From the point of view of algorithms, it is not that interesting to know merely that every polygon has an inside and an outside; we would like to know what they are! A standard test for this is the ray casting algorithm: given a point x and a polygon P , draw an infinite ray (or half-line) starting at x . Then x is inside P if the ray crosses P an odd number of times, and outside P if the ray crosses P an even number of times.

The choice of ray doesn't matter, and in programming applications, it's often fine to take a horizontal or vertical ray, for simplicity. We do run into some technicalities if the ray passes through a corner of P or contains an entire line segment of P . It's possible to make sense of this situation: we say that if a segment of P has an endpoint on the ray, that segment crosses the ray if it lies to the left of the ray, but not if it lies to the right. But this is very technical, and it's easier to avoid it if possible.

We can get a more general test as a consequence of this algorithm:

Claim 21.5. *Let P be a polygon, and let points x and y be connected by a polygonal curve whose segments never start or end on P , and never intersect a segment of P in a sub-segment, only in a point.*

Then x and y are on the same side of P if and only if the polygonal curve connecting them crosses P an even number of times.

Proof. We can subdivide the $x - y$ polygonal curve into segments that each cross P at most once. If a segment crosses P once, its endpoints are on opposite sides: this is verified by the ray casting algorithm with a ray parallel to the segment, starting at either endpoint. So with every crossing, the polygonal curve switches to a different side of P .

Since there are only two sides, the polygonal curve ends on the same side as it started if and only if it switched sides an even number of times. \square

We can now easily test which points lie in the same face, but do not have an easy definition of the boundary of a face. To obtain it, we need a geometric understanding of boundary walks.

Given a plane embedding, pick a value $\varepsilon > 0$ much smaller than every distance between a line segment and an endpoint of another end segment. Then surround every line segment by a thin rectangular loop: if the line segment has length ℓ , draw a $2\varepsilon \times (\ell + 2\varepsilon)$ rectangle with the line segment at its center. Orient each of these loops counterclockwise.

To find a boundary walk containing a particular segment, start going around its thin rectangular loop on either of the long sides. If this trajectory hits the loop around a different line segment (which only happens near the end of the segment), switch to following that loop. Keep going like this until this trajectory returns where it started (as it must, because we only have finitely many segments).

The result is a polygon that never crosses the plane embedding. It also stays within distance ε of a walk in the planar graph: it can only switch from following one edge to following another a vertex they share, because it never gets close enough to another edge anywhere else. We say that this walk is a boundary walk in the graph, and the polygon is the corresponding boundary walk in the plane embedding. Although boundary walks do not always represent cycles, we still consider two boundary walks to be the same if one is a shifted and/or reversed version of another.

The following is immediate from the definition.

- Each edge either lies on one boundary walk twice, or on two boundary walks, since each thin rectangular loop has two long sides.
- Each boundary walk in the plane embedding is contained entirely in one face, and we say it is a boundary walk of that face; a face can have multiple boundary walks.

Let xy be an edge of a planar graph G , and let C be the polygonal curve representing xy in a plane embedding of G . Many times, such as in the proof of Euler's formula, we want to understand edge xy by comparing the plane embedding of G and the plane embedding of $G - xy$ obtained by erasing C . In that second plane embedding, C stays entirely in some face F , because C does not cross any other edge. The following claim is a more careful formal proof of Lemma 21.2.

Claim 21.6. *Either $F - C$ is a single face in the plane embedding of G , xy lies on its boundary walk twice, and xy is a bridge of G ; or $F - C$ is the union of two faces in the plane embedding of G , xy lies on each of their boundary walks once, and xy is not a bridge of G .*

Proof. Starting from either side of a segment of C , construct boundary walks using xy , letting B_1 and B_2 be the corresponding polygons in the plane embedding; possibly, $B_1 = B_2$. If $B_1 = B_2$, then follow this polygon from one side of C to a point 2ε away on the other side, and connect by cutting across C . The result is a polygon P that only crosses C , and only once. No edge of $G - xy$ crosses P , and x and y are on opposite sides (because C crosses P once), so there is no $x - y$ path in $G - xy$: xy is a bridge. If $B_1 \neq B_2$, then each boundary walk in G uses edge xy only once, so it also contains an $x - y$ path in $G - xy$, proving that xy is not a bridge.

Each point $z \in F - C$ can be joined to a point of C by a polygonal curve not crossing the plane embedding of $G - xy$; before it touches C , it must cross either B_1 or B_2 . So all points of $F - C$

are in the face containing B_1 or the face containing B_2 . (The exception is points within ε of C , which can reach B_1 or B_2 by taking a step of length at most ε away from C .) So there are at most 2 faces, and if $B_1 = B_2$, there is only one face.

Conversely, if $F - C$ is a single face in the plane embedding of G , then take a line segment of length 2ε from B_1 to B_2 that cuts across C ; its endpoints are both in $F - C$, so this line segment can be completed to a polygon P that does not cross the plane embedding of G again. As before, no edge of $G - xy$ crosses P , and x and y are on opposite sides (because C crosses P once), so there is no $x - y$ path in $G - xy$: xy is a bridge. \square

The machinery we've built so far is enough to prove Euler's formula. As a special case, consider a planar graph G with no cycles: a forest. With n vertices and k connected components, the forest has $n - k$ edges, by Theorem 10.1. Substituting this into Euler's formula, we learn that the forest only has one face.

With at least two faces, the picture changes:

Claim 21.7. *If a plane embedding of G has at least two faces, then some boundary walk of every face contains a cycle.*

Proof. Let F be any face; since it is not the only face, we can pick a point $p \in F$ and a point $q \notin F$. The line segment pq leaves F , so at some point it must cross an edge xy on a boundary walk of F to do so. Replace pq by a short sub-segment, if necessary, to ensure that pq only crosses xy once, and does not cross any other edge of the plane embedding.

In the plane embedding of $G - xy$, points p and q are part of the same face (since pq no longer crosses any edge). By Claim 21.6, the boundary walk of F using xy only uses it once: after going from x to y , it is a $y - x$ walk that does not use edge xy . That $y - x$ walk contains a $y - x$ path, which together with edge xy becomes a cycle. \square

So far, we've been using the geometry of polygons to understand plane embeddings, but we can do this in reverse, as well.

Claim 21.8. *If p and q are two points on a polygon, then a polygonal curve that joins p and q divides one side of the polygon into two connected regions.*

Proof. Here, we think of the polygon plus the curve joining p and q as a plane embedding of a cycle graph with an extra edge (whose endpoints are drawn at p and q). That extra edge falls under the second case of Claim 21.6, proving the claim. \square

Claim 21.9. *If p, q, r , and s are four points appearing on a polygon P in that order, then a polygonal curve joining p and r cannot be drawn on the same side of the polygon as a polygonal curve joining q and s without crossing it.*

Proof. Suppose such a diagram existed; then it would be an embedding of K_4 , with vertices placed at p, q, r , and s . Such an embedding does exist: K_4 is a planar graph. But it cannot have the shape it needs to have here!

By Euler's formula, there are 4 faces total ($4 - 6 + 4 = 2$). One face of the plane embedding is the side P containing neither of the added curves; it has length 4. We know less about the others, but by Claim 21.7 each contains a cycle, so has length at least 3. The sum of the face lengths is at least $4 + 3 + 3 + 3 = 13$, but it must also be 12 by the face length formula, which is a contradiction. Therefore the diagram we started with cannot exist. \square

There is one final claim that we will need to construct new plane embeddings of graphs:

- Going from a plane embedding of G to a plane embedding of $G \bullet e$ in Chapter 22;
- Finding the dual graph of a plane embedding in Chapter 23;
- Finding a plane embedding of the graph of adjacencies of a map in Chapter 24.

Claim 21.10. *Given a plane embedding with a face F , a point $p \in F$, and points q_1, q_2, \dots, q_k on the boundary of F , it is possible to draw polygonal curves from p to each of q_1, \dots, q_k that don't cross each other or the boundary of F (and therefore stay inside F).*

Proof. We will draw the polygonal curves one at a time. First, draw a polygonal curve from p to a point on the same boundary walk as q_1 (which is possible by definition of both points being in F). Then, follow that boundary walk until its closest approach to q_1 , and end with a straight line segment.


(It is useful to observe that that last line segment can never intersect any part of the plane embedding: it is entirely contained in one of the rectangular loops around a segment of an edge, which is too close to that segment to touch any other segments.)

Once this is done, we can think of the result as a bigger plane embedding where the point p and the polygonal curve we drew are part of the boundary of a face F' contained in F . We add another polygonal curve by taking a short step out to the point p' on the boundary walk nearest p , then drawing a polygonal curve from p' to the next q_i in the same manner as before.


In general, this process can divide F' into two faces, and it's not guaranteed that the remaining points among q_1, q_2, \dots, q_k are all on the boundary of the same one. However, if this happens, then p is on the boundary of both of them, because the polygonal curve we drew from p is on the boundary of both, and we can repeat on each of the faces formed separately. \square

21.7 Practice problems

1. Find a plane embedding of the *volcano graph*, shown below.




2. Consider the following planar multigraph:



Draw the following plane embeddings:

- One where the outer face has length 3, with vertices 2, 3, 5 and edge e_1 on its boundary, and the other faces have lengths 2, 4, and 5.
 - One where the outer face has length 4, with vertices 1, 2, 4, 5 and e_1 on its boundary, and the other faces have lengths 3, 3, and 4.
 - One where the outer face has length 2, with edges e_1 and e_2 on its boundary.
- Let G be a connected planar graph. Suppose that a plane embedding of G has r faces, all of length 4.
 - Use the face length formula to find the number of edges of G (in terms of r).
 - Use Euler's formula to find the number of vertices of G (in terms of r).
 - Can you find such a graph for $r = 8$? What about for $r = 10$?
 - Use an argument similar to the proof of Proposition 21.1 to show that the complete graph K_5 is not planar. Start with a Hamilton cycle in K_5 , then reason about where the remaining 5 edges can go.
 - Prove that all trees are planar graphs. Moreover, show that any tree has a plane embedding in which all the edges are straight lines.
- (Hint: use induction. You'll need to show that if x is a leaf of T , then a plane embedding of $T - x$ can be extended to a plane embedding of T .)
- Here is another deceptive problem whose real solution can be found with Euler's formula. (It is documented along with many deceptive patterns in "The Strong Law of Small Numbers" by Richard Guy [11].)

Put n points around a circle in random positions (not equally spaced) and draw all $\binom{n}{2}$ chords between those points; wiggle the n points around, if necessary, until no intersection point inside the circle lies on three or more chords.



In the examples above (where $n = 1, 2, 3, 4, 5$) the chords divide the circle into 1, 2, 4, 8, and 16 pieces, so you might be forgiven for thinking that the number for general n is 2^{n-1} ; it's not! Use Euler's formula to find the correct answer.

7. a) Let G be a planar graph and let C be a cycle of length l in G . In every plane embedding of G , the cycle is a closed curve, but not necessarily the boundary of a face: its interior might be divided into several faces F_1, F_2, \dots, F_k by edges that are not part of C .

Prove that in all such cases, the sum $\text{len}(F_1) + \text{len}(F_2) + \dots + \text{len}(F_k)$ has the same parity as l : both are even, or both are odd.

- b) Use part (a) to prove that if G has a plane embedding in which every face has even length, then G is bipartite.

22 Planarity testing

The purpose of this chapter

This chapter presents three different approaches by which we can (sometimes or always) determine whether a graph is planar.

The first is Theorem 22.2, which is also very useful as a property of planar graphs: in Chapter 24, you will see how. It is especially important if you are not yet a very experienced mathematician to remember that this is very much not an if-and-only-if condition; do not use it in the wrong direction!

The second is Kuratowski's theorem and Wagner's theorem, which I group together, because they are very similar. As a planarity test, looking for minors is more powerful than looking for subdivisions, but minors are harder to understand—when teaching this material, I have often skipped Wagner's theorem for that reason. I include it here because in Chapter 24, it will be useful for us to know about edge contractions.

The third is the method of overlap graphs developed by Tutte. This is not as commonly studied, but I don't know why, because it's great. Any other equally systematic approach to planarity testing is much harder to learn.

22.1 Counting edges in planar graphs

In the previous chapter, we proved two tools that help us count vertices, edges, and faces in a plane embedding: the face length formula (Lemma 21.3) and Euler's formula (Theorem 21.4). We will now use these to prove a limit on the number of edges that a planar graph with n vertices can have.

For this, we will have to restrict our attention to simple graphs only, even though both results above apply to multigraphs as well. Otherwise, there is no possible bound we can prove!

Question: Why can a planar multigraph with n vertices have arbitrarily many edges?

Answer: Starting from an arbitrary plane embedding, we can draw in any number of loops at a vertex without crossings—imagine a flower with the loops as the petals, and the vertex at the center. If there is more than one vertex, we can also replace an edge by any number of parallel copies without affecting planarity.

We can engage in some meta-reasoning and ask: if our argument is to work for simple planar graphs, but not for planar multigraphs, what must it be doing? What sort of situations can only appear in the plane embedding of a multigraph?

Well, once again, we're back to the way that we can draw loops and parallel edges in a plane embedding. These are not unusual: they are the way we draw loops and parallel edges by default in a diagram of a multigraph. From the point of view of the counting tools we have, though, what makes them special is that they are boundaries of a face of length 1 (in the case of a loop) or 2 (in the case of two parallel edges).

Question: Do loops and parallel edges have to be drawn as faces of length 1 or 2?

Answer: No: at least in the case of some graphs, it's possible to draw the graph so that some part of it is inside the loop, or between the two parallel edges. (An example appears in the first practice problem at the end of the previous chapter.)

Question: Are there any simple graphs with faces of those lengths?

Answer: Just one: a graph with two vertices and a single edge between them has a plane embedding in which the outer face has length 2.

The following short lemma is the property we use to distinguish simple graphs from multigraphs. (From now on until we conclude our discussion of the number of edges in a planar graph, I will assume all graphs are simple.)

Lemma 22.1. *In a plane embedding of a simple graph with at least 2 edges, every face has length at least 3.*

Proof. If the plane embedding has multiple faces, then each face needs to be separated from the other faces somehow; its boundary must contain a closed curve, which corresponds to a cycle in the graph. The graph is simple, so a cycle in it must contain at least 3 edges.

If the plane embedding has only one face, then every edge of the graph must contribute 2 to the length of that face. There are at least 2 edges, so the face has length at least 4. \square

When we feed Lemma 22.1 into the tools we have from the previous chapter and let it cook, we obtain the following theorem.

Theorem 22.2. *For all planar graphs with $n \geq 3$ vertices and m edges, $m \leq 3n - 6$.*

Proof. Choose a plane embedding of a graph with $n \geq 3$ vertices and m edges to work with for the rest of the proof. Let r be the number of faces it has; by Lemma 22.1, each face has length at least 3, so the sum of the lengths of all the faces is at least $3r$.

Question: What if Lemma 22.1 does not apply, because $m < 2$?

Answer: Then $m \leq 3n - 6$ automatically, because $3n - 6 \geq 3(3) - 6 = 3$.

By the face length formula, the sum of the lengths of all the faces is also equal to $2m$, giving us the inequality $2m \geq 3r$. We would like to combine this inequality with Euler's formula: $n - m + r - k = 1$. We do not know k , the number of connected components, but it's certainly at least 1, so we get a second inequality: $n - m + r \geq 2$.

Our goal in this proof is to establish a relationship between m and n (though it can be interesting to look at the other two pairs of variables, too), so we should eliminate r . We solve for r in the inequality $2m \geq 3r$ (getting $r \leq \frac{2}{3}m$) and in Euler's formula (getting $r \geq m - n + 2$). Putting them together, we get

$$m - n + 2 \leq \frac{2}{3}m$$

which we can rearrange to $\frac{1}{3}m \leq n - 2$, or $m \leq 3n - 6$. □

Question: What would we learn if we eliminated m instead of eliminating r ?

Answer: From $m \leq n + r - 2$ and $m \geq \frac{3}{2}r$, then $n + r - 2 \geq \frac{3}{2}r$, or $r \leq 2n - 4$. This tells us the maximum number of faces in an n -vertex planar graph.

This chapter is called “Planarity testing”, and Theorem 22.2 is our first planarity test: it lets us immediately conclude that some graphs are not planar! Here is an example:

Proposition 22.3. *The complete graph K_5 is not planar.*

Proof. We count: K_5 has $n = 5$ vertices and $m = 10$ edges. Since $3n - 6 = 9$, which is less than m , the conclusion of Theorem 22.2 does not hold. Therefore the hypotheses do not hold either, so K_5 is not a planar graph. □

Question: If an n -vertex graph has $3n - 6$ or fewer edges, can we conclude from Theorem 22.2 that it is planar?

Answer: No: the theorem gives a necessary condition for planarity, but not a sufficient one. For example, take K_5 and add 95 isolated vertices, and you'll get a 100-vertex graph with only 10 edges which is not planar.

22.2 Triangulations

Whenever we prove an inequality, a natural question to ask is: what can we say about the cases where equality holds? What kind of planar graphs have $m = 3n - 6$?

To draw conclusions about such graphs, we should look back at our proof, and look at every place where an inequality appeared. This is an extremely useful idea not just in graph theory, but in other areas of math!




Figure 22.1: Triangulations with 7 vertices

1. We wrote Euler's formula as an inequality: $n - m + r \geq 2$. If it were the case that $n - m + r > 2$, we would conclude that $m < 3n - 6$.

So if a planar graph with $n \geq 3$ vertices satisfies $m = 3n - 6$, then $n - m + r$ must be equal to 2, meaning that the graph is connected.

2. The inequality $m \leq 3n - 6$ combines Euler's formula with the inequality $2m \geq 3r$. If we had started with the strict inequality $2m > 3r$, we would have arrived at the strict inequality $m < 3n - 6$, instead.

So if a planar graph with $n \geq 3$ vertices satisfies $m = 3n - 6$, it must satisfy $2m = 3r$.

3. The inequality $2m \geq 3r$ itself comes from another inequality, but one we only stated in words: Lemma 22.1, which says that every face has length at least 3. (For all faces F , $\text{len}(F) \geq 3$.)

So if a planar graph with $n \geq 3$ vertices satisfies $m = 3n - 6$, then every face must have length exactly 3, and this must be true in every plane embedding of the planar graph.

Now that we've understood these extremal examples, we give them a name, based on the fact that all of their faces are (in some sense) triangles. We call a plane embedding of a connected graph in which all faces have length 3 a **triangulation**. For each n , there are many n -vertex triangulations; see Figure 22.1 for some examples when $n = 7$. (These are all the 7-vertex possibilities, up to isomorphism of the planar graph. I know this because I first looked up sequence A000109 in the OEIS [21] to confirm there are 5 of them, then searched the House of Graphs [4] to find out what they are.)

To be clear, "triangulation" is a term for the plane embedding, not for the planar graph. How do we refer to the planar graph, then? The corresponding notion is a *maximal planar graph*: a planar graph that stops being planar if any edge (that it does not already have) is added to it. But to avoid sneaking in a claim that has not been justified, we need the following proposition.

Proposition 22.4. *For a planar graph G with $n \geq 3$ vertices, the following are equivalent:*

- (i) G has $3n - 6$ edges.
- (ii) Every plane embedding of G is a triangulation.
- (iii) G is a maximal planar graph.

Proof. We have already proven that (i) \iff (ii). G has exactly $3n - 6$ edges if and only if every inequality in the proof of Theorem 22.2 is an equality: if and only if every face has length exactly 3. This has to happen regardless of the plane embedding we choose at the beginning of that proof, so it must be true for every plane embedding.

Question: Which other implication between (i), (ii), and (iii) is easiest to show?

Answer: The implication (i) \implies (iii) also follows from Theorem 22.2. If G has $3n - 6$ edges, and we add a new edge to G , then we get a graph with $3n - 5$ edges, so by the contrapositive of Theorem 22.2, the new graph is not planar.

The part of Proposition 22.4 that we still need to prove is that (iii) also implies (i) and (ii). In other words, there are no maximal planar graphs that “get stuck” before reaching $3n - 6$ edges. We will show that (iii) implies (ii), and we will do it by showing the contrapositive: if G has an embedding that’s not a triangulation, then G is not a maximal planar graph.

Our strategy has a short description: in a plane embedding of G that isn’t a triangulation, we find a face F with $\text{len}(F) \geq 4$, and then we add an edge between two vertices on the boundary of F , drawing it inside F . The reason this is not the entire proof is that we have to make sure the edges do not already exist inside F .

In order for that to even be a problem, F must not be the only face, which means that there is a cycle on the boundary of F separating it from the other faces. This is either a cycle of length 4 or more, or else a cycle of length 3 with some more vertices of F “inside” the cycle; in either case, there are at least 4 vertices on the boundary of F .

If there are 5 or more vertices on the boundary of F , then they cannot all be adjacent: G would then contain a copy of K_5 , which we know cannot be drawn in the plane. If there are only 4 vertices, and they are all adjacent, then the plane embedding of G would contain a plane embedding of K_4 in which F is a face. This is impossible: by (i) \implies (ii), every plane embedding of K_4 is a triangulation, and cannot contain a face like F of length more than 3.

Question: Does anything change if F is the outer face?

Answer: Not substantially; the outer face is outside the cycles on its boundary, rather than inside them, but nothing changes aside from that one word.

In all cases, at least two vertices x, y on the boundary of F are not already adjacent in G . If we add edge xy to G , the result is still planar, because drawing a curve from x to y inside F gives a plane embedding of $G + xy$. This completes the proof: it shows that G is not a maximal planar graph. \square

22.3 Girth and planarity

Before we go on to necessary and sufficient conditions for planarity, let’s try to squeeze a bit more power out of the approach used to prove Theorem 22.2, because counting edges is a simpler test than just about anything else we could try.

The most common scenario is when G is a bipartite planar graph. In this case, the boundary of a face cannot be a cycle of length 3: by Theorem 13.1, bipartite graphs cannot contain such cycles! This allows us to sharpen our inequality.

Question: What if the boundary of a face is not a cycle?

Answer: In general, the length of the face is the total length of its boundary walks, and in a bipartite graph, all closed walks have even length.

If the minimum length of a face is 4, then we replace the inequality $2m \geq 3r$ in the proof of Theorem 22.2 by the inequality $2m \geq 4r$. As before, when we want an upper bound on m , we may assume G is connected, which lets us apply Euler's formula: $n - m + r = 2$. Eliminating r and simplifying, we conclude:

Theorem 22.5. *For all planar bipartite graphs with $n \geq 3$ vertices and m edges, $m \leq 2n - 4$.*

To appreciate the power of this approach, we can go back to Proposition 21.1 from the previous chapter. At the time, we need a technical argument that looked closely at the geometry of a plane embedding in order to prove that $K_{3,3}$ is not planar. With Theorem 22.5, the proof is simple: $K_{3,3}$ is a bipartite graph with $n = 6$ vertices and $m = 9$ edges. $9 > 2 \cdot 6 - 4$, so $K_{3,3}$ is not planar.

This argument is just one case of an even more general approach. Define the *girth* of a graph G to be the length of the shortest cycle in G . (This is always at least 3, and in bipartite graphs it is always at least 4.) In a graph with no cycles at all (a forest), the girth is sometimes defined to be ∞ , but that will not matter in this chapter. In any case, we don't need a planarity test for forests: they are always planar.

Theorem 22.6. *Let G be a planar graph with at least one cycle.*

If G has $n \geq 3$ vertices, m edges, and girth g , then $m \leq \frac{g}{g-2}(n-2)$.

Proof. Fix a plane embedding of G . It will have at least two faces, because a cycle separates the plane into two regions. Conversely, when there is more than one face, every face has a cycle in its boundary. Therefore the girth g is also a lower bound on the length of the faces of the plane embedding.

Now we can continue as before. Let the plane embedding have r faces; then the sum of the lengths of the faces (which is $2m$, by the face length formula) is at least rg . Combining the inequality $2m \geq rg$ with the equation $n - m + r = 2$ from Euler's formula, we get

$$2 - n + m = r \leq \frac{2m}{g}.$$

This simplifies to $m \leq \frac{g}{g-2}(n-2)$, the inequality we wanted. □




Figure 22.2: Non-planar graphs based on K_5

22.4 Subdivisions

Thus far, the graphs $K_{3,3}$ and K_5 have not been planar because they have too many edges. It is easy, but not very exciting, to find more graphs that are not planar, simply because they have a copy of $K_{3,3}$ or K_5 inside them. Figure 22.2b shows an example of this. We've added several vertices and edges to K_5 ; the result is even connected. However, the extra vertices and edges aren't really contributing anything to the non-planarity; the only reason that the resulting graph is not planar is because it contains a copy of K_5 inside it.

But we can consider stranger things than just subgraphs. Take a look at the graph in Figure 22.2c, for example. This graph does not have K_5 as a subgraph: it has 5 vertices in all the same places as Figure 22.2a, but one of the edges is missing. In place of that long edge is a path through some entirely new vertices.

Question: Why is the graph in Figure 22.2c not planar?

Answer: If we could draw it in the plane, we could simply erase the dots on the long path, and get a drawing of K_5 .

We can generalize this construction. We say that to **subdivide** an edge xy in a graph G means to create a new vertex z and replace edge xy by edges xz and yz . (In a diagram, this corresponds to drawing a dot representing z in the middle of edge xy .)

A **subdivision** of a graph G is a graph that can be obtained from G by subdividing edges zero or more times. (We consider G itself to be a subdivision of G .)

Question: Motivated by Figure 22.2c and our argument for it, what is the relationship between subdivisions and planarity?

Answer: If H is a subdivision of G , then G is planar if and only if H is planar: subdividing edges does not change planarity.

There is a reason why we have focused on two specific graphs that are not planar: the graphs K_5 and $K_{3,3}$. That reason is the following theorem, proved in 1930 by Casimir Kuratowski [17]:

Theorem 22.7 (Kuratowski's theorem). *A graph G is planar if and only if it contains a subdivision of K_5 or $K_{3,3}$.*

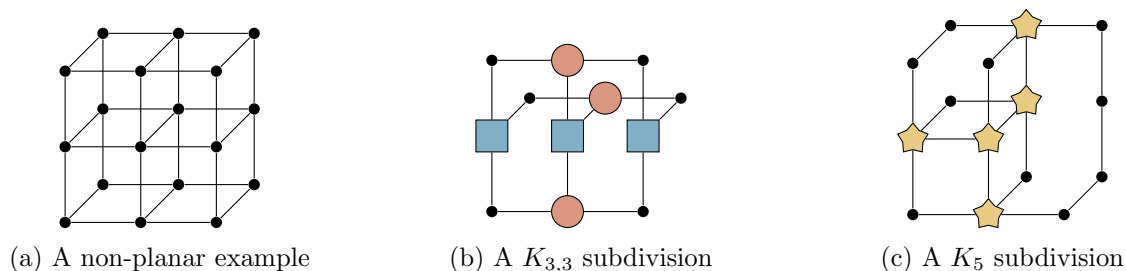


Figure 22.3: Showing that a graph is not planar using Kuratowski's theorem

Thus, subdivisions of K_5 and $K_{3,3}$ are the only reasons why a graph might fail to be planar. (Strictly speaking, I should write, “ G contains a copy of a subdivision of K_5 or $K_{3,3}$, but nobody is that strict about it, so I won't be, either.”)

Kuratowski's theorem is another example of a theorem that lets us concisely identify a reason why a problem in graph theory might have no solution. This is not necessarily related to how hard a problem is; it might be quite hard to find a plane embedding of a large graph, and it might also be hard to find a subdivision of K_5 or $K_{3,3}$ inside it. However, once you have a plane embedding, you can show it to anyone else, and instantly (or very quickly) convince them that a graph is planar. Working directly from the definition, there is not a lot you can show someone to quickly convince them that a graph is not planar.

That's exactly what Kuratowski's theorem provides. The subdivisions of K_5 and $K_{3,3}$ are obstacles to planarity. But Kuratowski's theorem mostly isn't about telling us that they're obstacles: we already knew that part. If we believe that K_5 and $K_{3,3}$ are not planar, then we can know that their subdivisions are not planar just by an argument like the one we used for Figure 22.2c. No, Kuratowski's theorem is about guarantees: it tells us that these obstacles are the only ones we need to worry about.

I will not give you a proof of Kuratowski's theorem. However, I will show you how we can look for a subdivision of K_5 or $K_{3,3}$ in a graph that is not too large. (If the graph is very large, then the task should be delegated to a computer, but it helps to have some understanding of what the computer could be trying to do.)

The example I will use is the graph in Figure 22.3a, which I will refer to as G for the rest of this section. A subdivision of $K_{3,3}$ inside G is shown in Figure 22.3b, and a subdivision of K_5 in Figure 22.3c. But how do we find them?

Question: Must both kinds of subdivisions exist in G , if it is not planar?

Answer: Not necessarily; Kuratowski's theorem only promises one of the two kinds of subdivisions. Even if both exist, we don't need to find both!

A first step might be to decide whether it is a subdivision of $K_{3,3}$ or a subdivision of K_5 that we want. If it exists, the subdivision of K_5 might be easier to look for, because K_5 has fewer vertices. However, K_5 has five vertices of degree 4, and this remains true for every subdivision of K_5 . If we're testing a graph G with fewer vertices of degree 4 or more, we can give up on finding a subdivision of K_5 , and focus on $K_{3,3}$. But G has enough degree-4 vertices (and even two degree-5 vertices) to find both.

In either case, we should begin by deciding which vertices will be the “key” vertices of the subdivision: the vertices of degree 3 or 4, as opposed to the vertices of degree 2 in the middle of subdivided edges. It is a good idea to pick central and high-degree vertices of the graph, to make it easier to find the paths later. A natural first choice is one of the two “center” vertices of G . (Call these vertices x and y , for future reference.)

If we did not know whether G is planar, we might also spend some time trying to find a plane embedding of G . This can also tell us something, even if we fail! For example, after trying for a while, you might be able to find a plane embedding of almost all of G , which is missing the edge xy between the two center vertices.

Question: What would such an embedding tell us?

Answer: Since $G - xy$ is planar, it does not contain a subdivision of $K_{3,3}$ or K_5 . Therefore whatever subdivision we hope to find in G must use edge xy somehow.

It is not, logically speaking, necessary for vertices x and y to be two of the key vertices of the subdivision. But it already seemed like a good idea to use these vertices before, so we might as well start with that theory.

It might take some trial and error to place the remaining key vertices. It is often a good idea to place as many of them close by as possible, so that they can be connected directly by edges; then, only a few long paths are necessary to draw the remaining connections.

It can help to switch between different ways of thinking about the graphs we’re subdividing—especially $K_{3,3}$, where we can either imagine the bipartition with three vertices on each side, or the hexagon with three chords. For a subdivision of K_5 , we can break down the process of finding it into stages:

1. Choose x and y to be our first two vertices.
2. Use edge xy as an edge of the subdivision, but also find three more longer $x - y$ paths, sharing no other vertices.
3. From each of those $x - y$ paths, choose a vertex to be a key vertex of the subdivision. Now, find a cycle through those three vertices, using no other the vertices used in previous stages.

A similar breakdown into stages could also work in other examples.

22.5 Graph minors

Another useful operation that preserves the planarity of a graph is edge contraction. If xy is an edge of graph G , the graph $G \bullet xy$ (also sometimes denoted G/xy) is the graph obtained from G by deleting vertices x and y and adding a new vertex z adjacent to every vertex which, in G , was adjacent to either x or y . The operation of going from G to $G \bullet xy$ is called **contracting** edge xy .




Figure 22.4: Examples of edge contractions and minors

In some applications, it makes sense to define $G \bullet xy$ as a multigraph. In that case, for every edge between a vertex w and either x or y in G , there is an edge between w and z in $G \bullet xy$; if there are multiple edges between x and y , and only one of them is contracted, the remaining edges become loops at z . In our case, loops and parallel edges are not relevant to track, so we will keep $G \bullet xy$ a simple graph.

Intuitively, given a plane embedding of G , we obtain a plane embedding of $G \bullet xy$ by first erasing everything within a small distance of edge xy (including the endpoints x and y). Within the erased region, draw vertex z , and change the trajectory of every edge formerly incident to x or y , so that instead it heads to z once it enters the erased region. Figure 22.4a shows an example of this intuitive idea, contracting an edge in a 3×4 grid graph.

Contracting edges is almost, but not quite, the reverse operation of subdividing edges. If we subdivide an edge xy (creating a new vertex z) and then contract edge xz (calling the combined vertex x) then we obtain the original graph again. However, contracting edges can do more than just undo edge subdivisions, when both endpoints of the contracted edge have degree 3 or more.

We say that a graph H is a *minor* of another graph G if it is a subgraph of a graph obtained from G by contracting edges zero or more times. If G is planar, then every minor of G is also planar, because contracting edges will not affect planarity, and neither will going from G to a subgraph.

Rather than think of a minor H as the result of a sequence of operations done to G , it can help to trace back where every vertex of H comes from.

Question: What is the simplest “origin story” of a vertex of H ?

Answer: It could have been a vertex of G to begin with.

Question: What is next simplest?

Answer: A vertex of H could have started as an edge of G .

Question: What else could have happened?

Answer: A vertex of H could be the result of contracting an edge whose endpoints were themselves the results of edge contraction(s). We could have an arbitrarily large set of vertices in G that all collapse down to one vertex in H .

There is only one restriction on the set of vertices that collapse down to one vertex in H . Such a set only shrinks due to contracting an edge between two of its vertices, so if we can get it down to a single vertex, it's because the set was originally a connected subgraph of G . So, we can identify the vertices of H with disjoint connected subgraphs of G ; if two vertices of H are adjacent, then the corresponding subgraphs of G must have at least one edge between them.

Figure 22.4b and Figure 22.4c show two examples of this. Well, to be precise, they show only the subgraphs which are to be contracted to vertices. In Figure 22.4b, the three vertical (blue) subgraphs become one side of $K_{3,3}$, and the three horizontal (red) subgraphs become the other side. To verify that we really do get a $K_{3,3}$ minor, we need to check that between each red vertex and each blue vertex, there is at least one edge. Verifying a K_5 minor takes less explanation: in Figure 22.4c, we need to check that there is an edge between every pair of the 5 circled regions.

Question: Why is finding a $K_{3,3}$ minor or a K_5 minor in a graph G interesting?

Answer: It shows that G is not planar: if G were planar, then all its minors would be planar, which is not true of $K_{3,3}$ or K_5 .

Klaus Wagner was the first to study graph minors in 1937 [24]. Among other results, he proved the analog of Kuratowski's theorem for graph minors:

Theorem 22.8 (Wagner's theorem). *A graph G is planar if and only if it does not have a minor isomorphic to $K_{3,3}$ or K_5 .*

As a characterization of planar graphs, Wagner's theorem follows from Kuratowski's theorem: if G contains a subdivision of $K_{3,3}$ or K_5 , then by contracting all the edges along paths in that subdivision, we can obtain a $K_{3,3}$ or K_5 minor. As far as we're concerned in this book, Wagner's more significant contribution is the concept of graph minors. Not all minors arise from subdivisions; in fact, it's possible to find a minor isomorphic to H in a graph which does not contain any subdivisions of H . So looking for a $K_{3,3}$ minor or a K_5 minor is easier than looking for a subdivision—once you're comfortable with the definition of a minor.

Outside the study of planar graphs, Wagner's work proved to be a greater influence on graph theory, because classifying graphs using their minors turned out to be a much more useful approach than classifying graphs using their subdivisions.




Figure 22.5: Showing that a graph is not planar using overlap graphs

22.6 Overlap graphs

A different approach to testing whether a graph is planar was first described in 1959 by Tutte [23]. It is less well-known than Kuratowski's theorem, but in my opinion, it is easier to use in small examples, and it can be the basis for efficient algorithms for planarity testing, as well. What's more, it is not just a way to prove that a graph is not planar; it can help find a plane embedding.

The idea is similar to the approach we took in the previous chapter to prove that $K_{3,3}$ is not planar. It begins by choosing a cycle in the graph we want to test for planarity. A Hamilton cycle works best, but of course Hamilton cycles can be hard to find, and we don't want to turn an easier problem into a harder problem! We can try the test with any cycle; however, the longer it is, the better.

As an example, I will show you how this approach works on the graph G in Figure 22.3. I have drawn it again with the vertices labeled in Figure 22.5a for two reasons: for ease of use, and also to point out a Hamilton cycle. In Figure 22.5b, that Hamilton cycle is drawn around a circle, and this picture is good to keep in mind for intuition.

In a hypothetical plane embedding of G , this cycle (and any other cycle) must appear as a closed loop of some sort. The rest of G must be drawn somewhere either inside that closed loop or outside it. In this example, "the rest of G " is just the edges which are not part of the cycle.

In general, given the graph G and a cycle C , let S be the set of vertices on C . We gather the vertices and edges of G which are not on C into *fragments*, defined as follows:

- Every edge of G whose endpoints are both in S is a fragment.
- Each connected component of $G - S$, together with the edges it has to S , is a fragment.

The motivation behind this definition is that a fragment is something that must either be drawn entirely inside C in a plane embedding of G , or else entirely outside it: it cannot cross C . For two different fragments, on the other hand, we can make different decisions. In our example, all fragments are edges, because C is a Hamilton cycle.

Question: Why shouldn't we have defined edges like $\{2, 5\}$ and $\{2, 17\}$, which share an endpoint, to be part of the one fragment?

Answer: Because we don't have to draw them on the same side of the cycle: to draw edge $\{2, 5\}$ inside the cycle and edge $\{2, 17\}$ outside it, no edges need to cross.

The decisions that we make for different fragments are not entirely independent. For example, edge $\{1, 6\}$ and edge $\{2, 17\}$ would intersect if we drew them both inside the Hamilton cycle, or both outside. To keep track of this information, we define a new graph, called the *overlap graph* of the cycle C . Its vertices are fragments, and they are adjacent if they "overlap": if they cannot be drawn on the same side of C .

A practical definition of the overlap graph in our case is that it has a vertex for each edge not part of C ; two vertices are adjacent if, in Figure 22.5b, they cross. This remains a perfectly good visual rule in general, but we'd like a combinatorial rule, because a computer cannot look at a drawing and see if two edges cross. The combinatorial rule for when two fragments overlap is this:

- There are four vertices w, x, y, z appearing in that cyclic order around the cycle C .
- One fragment includes an edge to vertex w and an edge to vertex y ; this may be satisfied by the fragment being edge wy .
- The other fragment includes an edge to vertex x and an edge to vertex z ; this may be satisfied by the fragment being edge xz .

The overlap graph in our example is shown in Figure 22.5c.

Guided by the overlap graph, we must choose for each fragment whether to draw it inside C or outside C . The rule is that if two fragment are adjacent in the overlap graph, then one of them must be inside C , and the other must be outside C .

Question: In terms of the overlap graph, what are we trying to find?

Answer: A 2-coloring, or bipartition! The fragments we decide to draw inside C are one side of the bipartition, and the fragments we decide to draw outside C are the other side.

We know from Chapter 13 that we can efficiently test if the overlap graph is bipartite, and Theorem 13.1 tells us that the overlap graph is either bipartite, or contains a cycle of odd length. That cycle is our proof that the graph cannot be drawn in the plane!

Question: Is there a cycle of odd length in the overlap graph in Figure 22.5c?

Answer: Yes: for example, the fragments $\{1, 6\}$, $\{4, 9\}$, and $\{5, 14\}$ form such a cycle.

Question: If the overlap graph were bipartite, would that be enough to conclude that G is planar?

Answer: Not in general: it's also possible that a single complicated fragment cannot be drawn together with C without crossing, regardless of what other fragments are doing. The overlap graph will not detect this.

However, in an example like this one, where C is a Hamilton cycle, deciding which fragments are inside and which fragments are outside C is all that there is to finding a plane embedding. A bipartition of the overlap graph tells us exactly how to draw G .

If the overlap graph is not bipartite, then it can also guide us to finding a subdivision of K_5 or $K_{3,3}$. The very simplest case is what we found in this example: three edges $\{1, 6\}$, $\{4, 9\}$, and $\{5, 14\}$ which all overlap. In that case, the cycle C together with those three edges is a subdivision of $K_{3,3}$.

In more complicated cases, more work needs to be done, but the overlap graph can still simplify the search. Take the cycle C , and the fragments forming an odd cycle in the overlap graph: just this portion of the graph alone is not planar, so it is all that is necessary to find a subdivision of K_5 or $K_{3,3}$.


22.7 Practice problems

1. The two graphs below were used as an example in the previous chapter.



One of these is planar, and the other one is not.

- a) Identify the planar graph, and draw a plane embedding.
 - b) Using one of the tools in this chapter, prove that the other graph is not planar.
2. The five connected 3-regular graphs with 8 vertices are all shown below. Determine which of them are planar, and which are not.



3. Prove that the Petersen graph is not planar in four different ways:
 - a) By using Theorem 22.6.
 - b) By finding a subdivision isomorphic to $K_{3,3}$.
 - c) By finding a minor isomorphic to K_5 .

- d) By finding a long cycle in the Petersen graph for which the overlap graph is not bipartite.
- 4. What is the maximum number of edges in an n -vertex planar graph if we know it has a plane embedding with two faces of length 6?
- 5. a) Let G be a connected graph with n vertices and $n + 2$ edges. Prove that G is planar.
b) Let G be a graph with n vertices and $n + 3$ edges obtained by starting with the cycle graph C_n and adding 3 more edges.

When is G planar, and when is G not planar?

- 6. The graph in Figure 22.3a is a $1 \times 2 \times 2$ three-dimensional grid graph. In general, the $a \times b \times c$ grid graph has vertices which are 3-dimensional points (x_1, x_2, x_3) with $x_1 \in \{1, 2, \dots, a\}$, $x_2 \in \{1, 2, \dots, b\}$, and $x_3 \in \{1, 2, \dots, c\}$; two vertices are adjacent if they are at distance 1 from each other. Thinking of Figure 22.3a as a 3-dimensional drawing can give you an idea of what an $a \times b \times c$ grid graph looks like.

For which a , b , and c is the $a \times b \times c$ grid graph planar?

- 7. (Putnam 2007) Let a triangulated polygon be a plane embedding in which every face except the outer face must have length 3. Prove that there is a function $f(n)$ such that if the outer face of a triangulated polygon has length n , and every vertex not on the boundary of the outer face has degree at least 6, then the triangulated polygon has at most $f(n)$ faces.

23 Polyhedra

The purpose of this chapter

To a graph theorist, Euler’s formula is a theorem about planar graphs. To almost every other mathematician, it is a theorem about three-dimensional solids. In this chapter, we’ll see the connection, and put graph theory to work in understanding 3D geometry.

In the middle of this chapter is a section on dual graphs, which I’ve included here because the duality between Platonic solids is a particularly striking instance of dual graphs at work.

The discharging method used to prove Theorem 23.3 is a common way to use Euler’s formula in proofs; it is also not a proof strategy that’s easy to come up with on your own. So I think it’s a particularly valuable proof to study, in case you ever encounter a problem in which this method can be used.

23.1 The Platonic solids

Regular polygons in the plane are abundant. For every $n \geq 3$, it is possible to draw a polygon with n equal sides and n equal angles. The sides can be however long we want them to be, but the angles must all have measure $\frac{n-2}{n} \cdot 180^\circ$ (or $\frac{n-2}{n} \cdot \pi$, in radians). The reason for this is that we can divide a regular n -sided polygon (or n -gon) into $n - 2$ triangles by drawing lines between its corners: two ways to do this are shown in Figure 23.1. In a triangle, the sum of angles is always 180° ; adding up the angles of all triangles gives $(n - 2) \cdot 180^\circ$, but this is also equal to the sum of all n angles of the regular n -gon.

Things change when we go up to 3 dimensions. The 3-dimensional version of a polygon is called a **polyhedron**. (The plural is “polyhedra”.) This is a shape with flat polygonal faces which come together at their sides and at their corners. If we want a polyhedron to be as regular as possible, then we can ask for the faces to be congruent regular polygons, with an equal number of them meeting at every corner.




Figure 23.1: Dividing regular polygons into triangles




Figure 23.2: The Platonic solids

This much regularity is a lot to ask for! In fact, there are only five possibilities once we ask for this much, shown in Figure 23.2. They are known as the *Platonic solids*, named after their description in Plato’s *Timaeus*, where Plato connects four of them to the four classical elements. They were analyzed more geometrically in Euclid’s *Elements*, and have been studied since then both mathematically and mystically [9].

Question: How many sides do the Platonic solids have?

Answer: Going from left to right in Figure 23.2: 4, 6, 8, 12, and 20. The Greek names of the solids refer to the number of sides: “tetra-” is a prefix that means 4, “octa-” means 8, and so on.

We will try to understand why there are just five of these solids. For this, it is necessary to connect the problem to graph theory somehow. You can probably already guess how we do it, by looking at the diagrams in Figure 23.2, and maybe recognizing the cube (Figure 23.2b) and dodecahedron (Figure 23.2d) from previous chapters. Each Platonic solid has an associated graph called its *skeleton graph*. Its vertices are the corners of the polyhedron, and its edges are the line segments where two of the polygonal faces meet.

(Actually, these objects are also called “vertices” and “edges” by geometers studying polyhedra. This is not a coincidence: graph theory gets its terminology from geometry, and not the other way around!)

However, the skeleton graphs of the Platonic solids are not just any kind of graph. They are planar graphs, and they have a plane embedding in which the faces (the polygonal sides of the polyhedron) become the faces of the plane embedding. This allows us to use the theory of planar graphs we have developed in the previous two chapters to solve a geometric problem.

There are two ways to get an intuition for how a polyhedron can be turned into a planar graph. One way is to imagine the polyhedron to be made of rubber; then, inflate the polyhedron until it becomes spherical, bending the edges into arcs on the sphere. (You should imagine the edges to be painted on, so that they do not get completely forgotten in this process.) Then, poke a hole in the rubber sphere, and stretch it out until it is flat; the drawing of the polyhedron on the sphere turns into a drawing in the plane.

Question: Is this possible for any polyhedron you can imagine?

Answer: No: some polyhedra have “holes” in them, and will inflate to a shape that is not a sphere. All five of the Platonic solids do inflate to spheres.




Figure 23.3: Going from a polyhedron to a planar graph

For some reason, mathematicians often feel that arguments involving imaginary rubber spheres are insufficiently rigorous. So here is another approach that is both visual and can be made rigorous. Hold the polyhedron above a horizontal plane, oriented so that a face on top is parallel to this plane. Place a bright light slightly above that face, close enough so that an observer at the bright light would see the top face and nothing else. Then the vertices and edges of the polyhedron will cast a shadow onto the horizontal plane, and that shadow will exactly be a plane embedding of the skeleton graph. An example is shown in Figure 23.3a.

The reason I say this can be made rigorous is that the projection via bright light can be described mathematically: for example, if the bright light is at point $(0, 0, 1)$ in \mathbb{R}^3 , and the horizontal plane is the plane $z = 0$, then the shadow of a point (x, y, z) has coordinates $(\frac{x}{1-z}, \frac{y}{1-z}, 0)$. Geometrically, the bright light, a point on the polyhedron, and its shadow are collinear.

We still need to take care to make sure that no edges cross in the shadow. A sufficient condition for this is to start with a *convex polyhedron*: one with the property that if two points A, B are contained in the polyhedron, then so is the entire line segment \overline{AB} . When a ray of light shines on a convex polyhedron, it always enters the interior of the polyhedron once (through the top face) and exits once. For two edges to cross in the shadow, the ray of light pointing at that crossing would need to enter and exit multiple times: once in the top face, and once for each edge that intersects at the crossing.

All of the polyhedra we consider in this chapter will be convex, and so all of them can be described as planar graphs—not just the Platonic solids. This correspondence can be taken even further, though. In general, a graph has a plane embedding exactly when it has a spherical embedding: when it can be drawn on the surface of a sphere without crossing. The reason is that we can take a spherical embedding and turn it into a plane embedding in exactly the same way that we’ve just done it for a polyhedron. Going the other way, a plane embedding can be drawn on the sphere just by using a small portion of the sphere that looks basically flat. We do this (with a certain very large sphere) every time we draw a picture in the dirt with a stick.

23.2 Classifying the Platonic solids

In two dimensions, there are infinitely many regular polygons. So why are there only five Platonic solids in three dimensions? This is, in part, something we can prove from Euler’s formula, with a quibble I will mention at the end of this section.

We can describe a Platonic solid by a pair (p, q) where every face has p sides, and q faces meet at every vertex. Geometrically, we must have $p \geq 3$ and $q \geq 3$.

Question: Why $p \geq 3$?

Answer: A polygon cannot have fewer than 3 sides.

Question: Why $q \geq 3$?

Answer: If we try to have only two polygons meet at a vertex, they end up lying flat against each other. I suppose I can't stop you from declaring that gluing two regular n -gons back-to-back is a Platonic solid with 2 faces at every vertex, but Plato wouldn't have been on board with this.

We can further narrow down the options for p and q by using the properties of a planar graph. (The planar graph must be connected, because a Platonic solid should be connected: two cubes floating in space next to each other are not a Platonic solid.)

Theorem 23.1. *There are only five possibilities for the pair (p, q) in a Platonic solid.*

Proof. We can write down two equations for n (the number of vertices), m (the number of edges), and r (the number of faces) in terms of p and q .

- The graph is a q -regular graph, so by the handshake lemma (Lemma 4.1), $nq = 2m$.
- Every face has length p , so by the face length formula (Lemma 21.3), $rp = 2m$.

We also have Euler's formula (Theorem 21.4): $n - m + r = 2$. Replacing n by $\frac{2m}{q}$ and r by $\frac{2m}{p}$, we get

$$\frac{2m}{q} - m + \frac{2m}{p} = 2 \implies \frac{1}{q} - \frac{1}{2} + \frac{1}{p} = \frac{1}{m}.$$

From here, the constraint that lets us narrow down the pairs (p, q) is that $\frac{1}{m} > 0$. Therefore $\frac{1}{q} - \frac{1}{2} + \frac{1}{p} > 0$, or $\frac{1}{p} + \frac{1}{q} > \frac{1}{2}$.

How can we get a total bigger than $\frac{1}{2}$ here? Let's do casework on p :

- If $p = 3$ (every face is a triangle) then $\frac{1}{q} > \frac{1}{2} - \frac{1}{p} = \frac{1}{6}$, so $q < 6$.
We can have $q = 3$ (three triangles meet at every vertex), giving us the tetrahedron.
We can have $q = 4$ (four triangles meet at every vertex), giving us the octahedron.
We can have $q = 5$ (five triangles meet at every vertex), giving us the icosahedron.
- If $p = 4$ (every face is a square) then $\frac{1}{q} > \frac{1}{2} - \frac{1}{p} = \frac{1}{4}$, so $q < 4$.
We can have $q = 3$ (three squares meet at every vertex), giving us the cube.
- If $p = 5$ (every face is a pentagon) then $\frac{1}{q} > \frac{1}{2} - \frac{1}{p} = 0.3$, so $q < \frac{1}{0.3} = 3\frac{1}{3}$.
We can have $q = 3$ (three pentagons meet at every vertex), giving us the dodecahedron.

These are the only possibilities: if $p \geq 6$, then not even $q = 3$ is small enough for the inequality $\frac{1}{p} + \frac{1}{q} > \frac{1}{2}$ to hold. \square

In each of these cases, once we know p and q , we can solve for n , m , and r .

Question: What must n , m , and r be if $(p, q) = (3, 4)$?

Answer: Starting from Euler's formula $n - m + r = 2$, we can substitute $n = \frac{2m}{q} = \frac{1}{2}m$ and $r = \frac{2m}{p} = \frac{2}{3}m$ to get $\frac{1}{2}m - m + \frac{2}{3}m = 2$, or $\frac{1}{6}m = 2$. Therefore $m = 12$, and now we can back-substitute to get $n = \frac{1}{2}m = 6$ and $r = \frac{2}{3}m = 8$. This is the octahedron: it has $n = 6$ vertices, $r = 8$ faces, and $m = 12$ edges.

In general, the equation $\frac{1}{q} - \frac{1}{2} + \frac{1}{p} = \frac{1}{m}$, which can be rearranged to get $m = \frac{1}{1/p+1/q-1/2}$. Then $n = \frac{2m}{q}$ and $r = \frac{2m}{p}$ tells us the number of vertices and the number of faces. Here is the complete table (where $\langle q \rangle \times n$ stands for the sequence q, q, \dots, q of length n):

	Tetrahedron	Cube	Octahedron	Dodecahedron	Icosahedron
Number of vertices	4	8	6	20	12
Degree sequence	$\langle 3 \rangle \times 4$	$\langle 3 \rangle \times 8$	$\langle 4 \rangle \times 6$	$\langle 5 \rangle \times 12$	$\langle 3 \rangle \times 20$
Number of edges	6	12	12	30	30
Number of faces	4	6	8	12	20
Face types	$\triangle \times 4$	$\square \times 6$	$\triangle \times 8$	$\diamond \times 12$	$\triangle \times 20$

I promised to mention a quibble I have with calling this a complete classification of the Platonic solids. Well, first of all, it is only a classification of the skeleton graphs of those solids: we have not (and will not) engage with the 3D geometry. Even so, one thing is missing.

Question: How could there, conceivably, be another regular planar graph where every face in a plane embedding has the same length, other than the ones in Figure 23.2?

Answer: It's possible that there are several different non-isomorphic graphs corresponding to a single pair (p, q) .

So is there a second icosahedron where the faces attach differently? There is not, but that takes some effort to prove. In principle, it's a finite problem: up to isomorphism, there are only finitely many graphs with 20 or fewer vertices, and we could simply check them all and verify that none of them have the regularity properties we asked for.

It would be nice to have a more elegant argument, though. For the smaller Platonic solids, this is achievable. For example, the pair $(p, q) = (3, 3)$ must correspond to a 3-regular, 4-vertex graph, and there is only one such graph: the complete graph K_4 . Slightly more complicated, but similar arguments work for $(p, q) = (3, 4)$ and $(p, q) = (4, 3)$; since I know it's possible to handle these two cases in a slick way, I will leave it for you to discover in the exercises at the end of this chapter. I am not aware of any arguments for $(p, q) = (3, 5)$ and $(p, q) = (5, 3)$ that do not require a substantial amount of suffering, so I will not make you deal with those cases.

23.3 Dual graphs

If you look at the table counting vertices, edges, and faces in the Platonic solids, you may notice an interesting pattern: the five solids can be divided into two and a half pairs for which these counts are related. The triple (n, m, r) counting the vertices, edges, and faces is $(8, 12, 6)$ for the cube and it is $(6, 12, 8)$ (the reverse) for the octahedron; it is $(20, 30, 12)$ for the dodecahedron and $(12, 30, 20)$ (the reverse) for the icosahedron.

Question: Why “two and a half” pairs?

Answer: The tetrahedron can be paired with itself, because it has the same number of vertices and faces.

There is another correspondence between vertices and faces that you may have noticed before. The face length formula for plane embeddings is very similar to the handshake lemma: if F_1, F_2, \dots, F_r are the faces and x_1, x_2, \dots, x_n are the vertices, then

$$\sum_{i=1}^r \text{len}(F_i) = 2m = \sum_{i=1}^n \deg(x_i).$$

Is this a coincidence?

Mathematicians should always be on the lookout for such “coincidences”, because it often turns out that they reveal a deeper idea. In this case, it leads to the definition of dual graphs.

The **dual graph** of a plane embedding is not really a graph, but a multigraph whose vertices are the faces of the plane embedding; for every edge in the plane embedding, there is an edge in the dual graph between the faces that meet there. For example, Figure 23.4d shows the dual graph of the plane embedding in Figure 23.4a.

Question: The definition mentions that the dual graph is actually a multigraph. When does it have loops?

Answer: When the original plane embedding has an edge with the same face on both sides. A vertex of degree 1 guarantees that this will happen, though it is not the only way.

Question: What about parallel edges?

Answer: The dual graph has parallel edges when two faces border each other along multiple edges. A vertex of degree 2 guarantees that this will happen, though it is not the only way.

The mere existence of the dual graph, carefully defined, is enough to derive the face length formula as a consequence of the handshake lemma. For relationships like those between the Platonic solids to hold, something more has to happen: the dual graph must itself be a planar graph (or multigraph).




Figure 23.4: The dual graph of a plane embedding

Proposition 23.2. *The dual graph of a plane embedding is also planar.*

Proof. Figure 23.4 shows the process of carefully constructing the dual graph in such a way that we get a plane embedding of the dual graph at the same time.

Begin by drawing a dual vertex somewhere in the interior of each face, and marking a crossing point somewhere in the middle of each edge. This is shown in Figure 23.4b.

Next, to draw a dual edge between the dual vertices inside faces F_i and F_j , we draw a curve from the dual vertex inside F_i , to the crossing point on the edge that F_i and F_j share, to the dual vertex inside F_j .

We still need to be careful to avoid crossings, but the setup means we need to be less careful. In order to know that this construction can always be carried out, we only need to know one thing: inside a face F_i , we can draw non-intersecting curves from the dual vertex to all the crossing points on the boundary of F_i . This is a “local” geometric claim that doesn’t require us to consider the plane embedding as a whole. \square

There are several properties that the dual graph isn’t required to have, but that it will have in sufficiently nice cases. For example, strictly speaking, it is incorrect to refer to the “dual graph of G ”, where G is a planar graph. The dual graph is defined based on a plane embedding of G , not based on G itself. Sometimes we can get legitimately different dual graphs by choosing a different plane embedding, though the graph in Figure 23.4a and the skeleton graphs of the Platonic solids do not allow this.

Question: From the plane embeddings we’ve seen so far, can you give an example where two different ones are guaranteed to give non-isomorphic dual graphs?

Answer: In Chapter 21, Figure 21.5a and Figure 21.5b showed two plane embeddings of the same graph where the faces had different lengths. The dual graphs we get from these embeddings will not be isomorphic, because their vertices will have different degrees.

In Figure 23.4c, only the color indicates which edges are edges of the original plane embedding, and which edges are edges of the dual graph. Here, the dual relationship holds in both directions: each vertex of the original plane embedding lies in a face of (the plane embedding of) the dual graph. In such a scenario, taking the dual graph twice brings us back to where we started,

up to isomorphism. However, it is not guaranteed to happen; it is even possible for a plane embedding to have more vertices than the dual graph has faces.

Given a connected plane embedding with n vertices, m edges, and r faces, however, the dual graph will always have n faces (as well as m edges and r vertices). The number of edges and vertices in the dual graph follows from the definition, but the number of faces follows from Euler's formula. Applied to the original plane embedding, it tells us that $n - m + r = 2$. Meanwhile, if we suppose that a plane embedding of the dual graph has n' faces, then $r - m + n' = 2$; together, these two equations imply that $n = n'$.

Question: Must the dual graph of a plane embedding always be connected?

Answer: Yes: given any two faces F and F' (which are vertices of the dual graph), draw any curve from the inside of F to the inside of F' , only making sure it does not pass through any vertex. The faces that the curve passes through form an $F - F'$ walk in the dual graph.

Finally, there is more to the story of duality in the case of Platonic solids (and some other polyhedra). For the skeleton graph of a polyhedron, we can construct a dual graph in a way that reflects the geometry of the polyhedron, by doing the following:

1. For the dual vertex corresponding to a face F of the polyhedron, draw a point in the geometric center of face F .
2. For the dual edge connecting adjacent faces F and F' , draw a line segment between the two points in the centers of F and F' .

Because the dual vertices are adjacent exactly when the corresponding faces are adjacent, this is still a geometric realization of the dual graph. (It is not a plane embedding because it's all happening in three dimensions.)

This kind of dual turns the cube into the octahedron (and vice versa) geometrically, not just graph-theoretically; the same is true for the icosahedron and dodecahedron!

Question: In general, though, this construction is not guaranteed to turn a polyhedron into another polyhedron. What could go wrong?

Answer: The resulting points and line segments do not necessarily form faces that lie flat! If the dual graph has a face of length 4 or more, then the points on the boundary of that face might not end up lying on a single plane.

23.4 Archimedean solids

The definition of a Platonic solid is the most restrictive generalization of a regular polygon to three dimensions. A slightly less restrictive, and still very interesting, definition is that of an *Archimedean solid*.




Figure 23.5: Truncating a tetrahedron

These are convex polyhedra whose faces are all regular polygons, and whose vertices are all symmetric to each other (that is, for any two vertices, there is some rotation or reflection of the polyhedron that can move one to the other). Notably missing from this definition is any kind of symmetry between faces: in an Archimedean solid, the faces do not all have to be the same!

The truncated tetrahedron is one small example of an Archimedean solid. Geometrically, it is obtained as follows: start with a tetrahedron, and cut off each vertex a third of the way along its edge, as shown in the picture below. The truncation process is shown in Figure 23.5.

The truncated tetrahedron has two types of faces: four hexagons (left over from the original faces of the tetrahedron) and four triangles (from where the cuts were made).

As with the Platonic solids, we can at the very least determine the global face, vertex, and edge counts from a local description of what is happening at every vertex. Let's see how, using the truncated tetrahedron as an example. (Imagine that we don't have the diagram in Figure 23.5b to use as a reference.)

Every vertex of the truncated tetrahedron is a third of the way along some edge of the tetrahedron. The two faces that meet there are two hexagons (from the two faces of the tetrahedron that met along that edge) and one triangle (from the cut that created that vertex). This is all the "local information" we will need. The variables we will solve for are:

1. m , the number of edges.
2. n , the number of vertices.
3. r_3 , the number of 3-sided faces (triangles).
4. r_6 , the number of 6-sided faces (hexagons).

We will need four equations, because there are four variables. Euler's formula is one of them: it tells us that $n - m + (r_3 + r_6) = 2$. It is tempting to use the face length formula, which tells us that $3r_3 + 6r_6 = 2m$, but this is less convenient because it involves three variables; instead, we take the handshake lemma, which tells us that $3n = 2m$. For the other two equations, we use the local information about what happens at each vertex.

Question: If there are n vertices, and each vertex the corner of one of the r_3 triangles, what is r_3 in terms of n ?

Answer: There is a 3-to-1 correspondence between vertices and triangles: each vertex has 1 triangle, but each triangle has 3 vertices. So $n = 3r_3$.

Question: What about the relationship between n and r_6 ?

Answer: Here, the correspondence is 6-to-2: each vertex has 2 hexagons that meet there, and each hexagon has 6 vertices at its corners. So $2n = 6r_6$.

In general, such equations are determined by two quantities: the number of sides each type of face has, and the number of faces of that type that meet at each vertex.

Now we can write Euler's formula solely in terms of n , by replacing each variable by a multiple of n : since $m = \frac{3}{2}n$ and $r_3 = r_6 = \frac{1}{3}n$, Euler's formula turns into

$$n - \frac{3}{2}n + \frac{1}{3}n + \frac{1}{3}n = 2.$$

When we simplify, we get $n = 12$. Therefore $m = \frac{3}{2}n = 18$, $r_3 = \frac{1}{3}n = 4$, and $r_6 = \frac{1}{3}n = 4$: exactly the parameters of a truncated tetrahedron!

There is a way to make this process more systematic, while also generalizing it to be able to deal with even less regular polyhedra.

To do so, we define the *angle defect* at a corner of a polyhedron to be 2π minus the sum of the angles of the polygons meeting at that corner. (We'll work in radians from now on; in degrees, we'd take 360° instead of 2π .) Since the angle defect would always be 0 if the corner were flat, this is a measure of how much the polyhedron "bends" at a corner.

For a convex polyhedron (or any polyhedron for which Euler's formula holds), no matter how many vertices there are, the total amount of "bend" must be the same. This was originally shown by René Descartes, the inventor of coordinate geometry [7]:

Theorem 23.3. *In any convex polyhedron, the sum of all angle defects is 4π .*

Proof. This could be done by solving a system of equations, but there is a more elegant proof by a strategy called the "discharging method". We take the skeleton graph of the polyhedron, and put a "charge" of $+2\pi$ on each vertex, $+2\pi$ on each face, and -2π on each edge. The total charge on the graph is $2\pi n - 2\pi m + 2\pi r$, and we've chosen our initial charges so that the total charge would simplify to 4π by Euler's formula.

In physics, positive and negative electric charges cancel. Probably. I'm not a physicist. In this proof, we will move around the charges we've placed on the polyhedron to cancel them, while not changing the overall sum. First, from each face, we move $+\pi$ charge on to each of its edges. This leaves each edge at charge 0; it started at -2π , but gained $+\pi$ from each of the two faces it borders. However, each face has now gone into the negatives: a face of length l now has charge $-(l - 2)\pi$.

In an l -sided polygon, the sum of angles is $(l - 2)\pi$ (as observed at the beginning of this chapter for regular polygons). So we can bring each face up to zero charge with a second transformation: for every corner of every face, if that corner makes an angle of θ on that face in the polyhedron, we move θ charge from the vertex at that corner to the face.


When we're done, the faces and edges all have charge 0, while the remaining charge at each vertex is exactly the angle defect. However, the sum of the charges has remained at 4π throughout, proving the formula. \square

Theorem 23.3 can be used to quickly count the vertices in any Archimedean solid. For example, in the truncated tetrahedron, a regular triangle and two regular hexagons meet at each vertex, forming one angle of measure $\frac{\pi}{3}$ and two angles of measure $\frac{2\pi}{3}$. This means that the angle defect at each vertex is $2\pi - \frac{\pi}{3} - 2(\frac{2\pi}{3}) = \frac{\pi}{3}$. The total angle defect is 4π , so there must be $\frac{4\pi}{\pi/3} = 12$ vertices.

23.5 Practice problems

1. Draw a plane embedding of the skeleton graph of the dodecahedron.
2. The skeleton graph of the icosahedron is pancyclic: it has a cycle of every length from 3 to 12. Verify this by finding a cycle of each length.
3. Determine the dual graph of any plane embedding of any n -vertex tree, up to isomorphism.
(This example goes to show that two planar graphs with isomorphic duals are not, themselves, necessarily isomorphic.)

4. Here are five different plane embeddings of a graph G :



For each plane embedding, draw the dual graph. Determine which of these graphs are isomorphic to each other, and which are not.

5. By going from a plane embedding to a spherical embedding and back to a plane embedding (using the projection technique in Figure 23.3), prove that for any face F of a plane embedding of G , there is a plane embedding of G where F is the outer face.
6. A uniform n -gonal prism is a prism with $n + 2$ faces: two regular n -gons on the top and bottom, and n squares around the sides.
 - a) Draw a plane embedding of the skeleton graph of a uniform n -gonal prism where n is some very big number—like 6. Explain how to draw such a plane embedding for any value of n .
 - b) Draw the dual graph of the plane embedding you drew in part (a). Describe the structure of this dual graph for arbitrary values of n .
 - c) Geometrically, what does the dual polyhedron of the n -gonal prism look like?
7. a) (AIME 2004) A convex polyhedron P has 26 vertices, 60 edges, 36 faces, 24 of which are triangular, and 12 of which are quadrilaterals. A space diagonal is a line segment connecting two non-adjacent vertices that do not belong to the same face. How many space diagonals does P have?
 - b) How many of the numbers given in part (a) are redundant information?
8. Here are few more questions about Archimedean solids.

- a) An icosidodecahedron is an Archimedean solid with 12 pentagonal faces (like a dodecahedron) and 20 triangular faces (like an icosahedron). How many vertices and edges does it have? How many faces of each type meet at each vertex?
 - b) A snub cube is an Archimedean solid with four triangles and one square meeting at every vertex. How many vertices and edges does it have, and how many faces of each type?
 - c) What about the truncated icosidodecahedron, in which a 4-sided face, a 6-sided face, and a 10-sided face meet at every vertex?
9. Suppose that G is an n -vertex planar multigraph such that (for at least one plane embedding of G) the dual graph is isomorphic to G . (We call such a multigraph self-dual.)
- a) How many edges must G have, in terms of n ?
 - b) Find an example of such a graph G for all $n \geq 2$.
10. In this problem you will prove that the standard octahedron and cube are the only possible Platonic solids with their parameters, at least from the point of view of graph theory.
- a) It follows from Theorem 23.1 that any Platonic solid with $(p, q) = (3, 4)$ is a 4-regular 6-vertex graph. Prove that there is only one such graph (up to isomorphism).
 - b) It follows from Theorem 23.1 that any Platonic solid with $(p, q) = (4, 3)$ is a 3-regular 8-vertex graph. Unfortunately, there are multiple such graphs. However, the graph must also have a plane embedding in which every face has length 4, and by a practice problem at the end of Chapter 21, it must be bipartite.
Prove that there is only one bipartite 3-regular 8-vertex graph (up to isomorphism).

24 Coloring maps

The purpose of this chapter

I have to write about map coloring at some point in this textbook; there's no way around it. The map coloring problem is not only the reason graph coloring was invented; it's also a big part of the story behind the invention of graph theory as a coherent discipline. That's also why there's more history in this chapter of the textbook than in most other chapters.

I don't usually spend a whole lecture on map coloring when teaching a course on graph theory, because there's not enough time. I think going up to just Theorem 24.4 is a reasonable minimum: it is a good way to see how the properties of planar graphs we already know can be applied, and it's a nice application of coloring graphs greedily. The proof of Theorem 24.5 I've included is a bit shorter than the usual one, specifically in case you've decided to spend a bit more time on the topic of map coloring, but not too much time. (As a side note, the edge contraction in the proof has a particularly nice representation if we are coloring a map, not a graph: then, we just erase the borders region x has with y_i and y_j .)

The last two sections cover two special topics that are less frequently seen in graph theory courses, but which I think are very interesting. (I think it's fascinating how—for the second time in our study of planar graphs!—finding a Hamilton cycle can help us solve a seemingly unrelated problem.) Heawood's empire coloring problem is a good example of how we can arrive at new mathematical questions by examining the assumptions in our simplified models.

In addition to the previous chapters in this book, this chapter heavily relies on Chapter 19 (naturally). On the other hand, the section on the use of Hamilton cycles in coloring maps will not ask you to know much more from Chapter 17 than the definition of a Hamilton cycle.

24.1 Coloring maps

In Chapter 1, we visited Switzerland; let us return there.

Figure 24.1 shows a map of the cantons of Switzerland. To make the borders between cantons easy to see, we give them different colors. We are willing to reuse colors, but we would like two cantons to have different colors when they share a border. (It would be fine for two cantons to have the same color if they only share a corner; this situation is rarely found in maps.) How many different colors do we need to color Switzerland—and is there a number of colors that would be enough to color any map, no matter how complicated?

Let the **graph of adjacencies** of a map be the graph whose vertices are the regions we must color, with an edge between vertices that share a border. Then what we are looking for is a (proper) coloring of the graph of adjacencies: we've arrived back at the graph coloring




Figure 24.1: A map of Switzerland (coordinates from [20])

problem studied in Chapter 19! Historically, in the middle of the 19th century, this was the first graph coloring problem studied. Graph theory had yet to be codified as a discipline at the time, though some problems were studied individually which are now the domain of graph theory. The book *Four Colors Suffice* by Robin J. Wilson [26] gives a history of the problem, and is also the source for most of my historical claims in this chapter.


Why are we only looking at the map coloring problem now? The reason is that for maps with reasonable, well-behaved regions, the graph of adjacencies is planar. We can show this by the same argument we used to prove Proposition 23.2 that the dual graph of a plane embedding is planar. In fact, we can think of the graph of adjacencies as the dual of a plane embedding whose edges are exactly the borders drawn in the map. (We add vertices to give these edges suitable endpoints; aside from a few edge cases, vertices are mostly only necessary where three borders meet.)

I said “maps with reasonable, well-behaved regions” because the real world is messy, and not all actual maps translate to planar graphs as neatly. In fact, Switzerland already provides an example of this: its cantons are not what I would call reasonable and well-behaved!


Question: What properties do the Swiss cantons have that the faces of a plane embedding shouldn’t?

Answer: They’re often not connected regions! Four cantons have multiple pieces separated by another canton: Obwalden (6), Fribourg (10), Solothurn (11), and Vaud (22).

In fact, we can prove that the graph of adjacencies between Swiss cantons is not a planar graph. This follows from Kuratowski’s theorem (Theorem 22.7). It’s enough to look at the cantons shown in Figure 24.2a: because canton 6 (Obwalden) has two pieces, there are a few more borders between these than would be possible in a reasonable map. The graph of adjacencies between these 6 cantons contains a subdivision of the complete graph K_5 , which is shown in



(a) A map of a few Swiss cantons



(b) A subdivision of K_5 in Switzerland

Figure 24.2: Why Switzerland is not planar

Figure 24.2b: 9 of the 10 edges between cantons $\{2, 3, 4, 6, 7\}$ are present directly, and instead of the edge $\{3, 4\}$ there are edges $\{3, 5\}$ and $\{5, 4\}$ through 5, an extra vertex. (Edge $\{5, 7\}$ is also present in the graph, but is not shown in Figure 24.2b because it is not part of the subdivision.) Since the subgraph in Figure 24.2b is not planar, the entire graph of adjacencies between Swiss cantons cannot be planar, either.

This is an extremely practical example of a misbehaving map; it is also possible to give another example that is extremely theoretical. Hud Hudson showed [13] that if we allow fractal regions with a finite area and infinite perimeter (considering them to share a border if both regions get arbitrarily close to that border), then even connected regions can touch in extremely non-planar ways. There is no limit to the number of colors that might be necessary to color such a fractal map.

From now on, we will assume that our maps are reasonable and well-behaved: that the borders between regions obey the same restrictions as edges in a planar graph (there are no fractal borders) and that the objects we are coloring are nothing more than the connected regions separated by those borders (there are no regions with multiple pieces). With these caveats, the problem of coloring maps is exactly equivalent to the problem of coloring planar graphs.

24.2 The four color theorem

What could stop us from coloring a map with k colors, for some integer k ? The most straightforward kind of obstacle is a set of k regions among which any two are adjacent: a k -vertex clique in the graph we are coloring. It's possible to draw a map with four regions that all touch; this is the simplest possible example of a map which requires four colors.

Question: Is it possible to draw a map with 5 regions that all touch?

Answer: No, because the graph of adjacencies would be K_5 , which is not a planar graph.

Question: Does this prove that four colors are enough to color any map?

Answer: No! It is possible to have a graph with no 5-vertex clique which is still not 4-colorable.

If you’ve been reading this book carefully and in order, Chapter 19 should already have prepared you for the idea that the chromatic number $\chi(G)$ and the clique number $\omega(G)$ are very different. Forgetting this is one of the most common mistakes people make when they first think about coloring planar graphs.

From a pedagogical point of view, I can’t help but think that it’s a shame that the following theorem really is true:

Theorem 24.1 (Four color theorem). *If G is a planar graph, then $\chi(G) \leq 4$.*

The road to proving this theorem has been a long one. The mathematical question of whether four colors are enough to color any map was first asked by Francis Guthrie in 1852. Later in the 19th century, it received a great deal of mathematical attention, which came with multiple incorrect solutions to the problem.

Some incorrect solutions were the sort of careless error that confuses the chromatic number with the clique number. I assume that even at the time, there were also many attempts to disprove the theorem, from people who drew a very complicated map and could not find a way to color it with four colors; there are still such attempts. (At the same time, there was very persuasive experimental evidence for the four color theorem: every real-world and imaginary map that was tried could in fact be colored with four colors.)

There were also more notable attempts. Later on in this chapter, I will mention two attempts at proving the four color theorem that were eventually found to be false, but contributed a lot to graph theory, even so. Alfred Kempe’s approach had a subtle error, but it was still enough to give a proof of the weaker bound in Theorem 24.5, and the ideas in Kempe’s proof have since been used to solve other problems about graph coloring. And I suspect that Peter Tait’s use of Hamilton cycles to try to color maps was one of the big reasons why 19th century mathematicians continued to study Hamilton cycles, before practical uses of them were found!

Eventually, a proof of the four color theorem was found, but even that was controversial. Kenneth Appel and Wolfgang Haken, in a long effort from 1972 to 1976, found a proof that relied on using a computer to verify 1936 different configurations [2, 3]. How can any number of finite cases be used to prove a theorem about the infinite variety of possible maps? The idea, simplified, is that any possible map would be guaranteed to contain one of the “reducible configurations” that Appel and Haken found. The configurations were not just substructures that could be colored: they could be replaced by smaller ones without hurting colorability, which allows for a proof by induction on the number of regions in the map.

A proof so long that it needed a computer to verify was controversial. Wouldn’t it be invalidated by a single bug in the program? This was the fear at first, but by now, it’s been nearly 50 years, and we can be a bit more confident, for a few reasons:

- A 1997 proof by Robertson, Sanders, Seymour, and Thomas, while still a computer-based proof, used fewer configurations and a simpler approach to reductions [19].
- In 2005, Georges Gonthier formalized the entire proof in the Coq proof assistant [10]. Though the proof is still checked by computer in this case, it does not rely on rules specific to map coloring, only on formal logic, so we can be more confident in the computer.
- It’s been nearly 50 years! If there were something incurably wrong, surely we’d have discovered it by now.

Even if we believe the computer, there is still a reason why we might want a short, human-readable proof (which still has not been found). Such a proof would certainly contain new ideas that we can apply to solve other, more difficult problems! This is not to say that the computer-based proofs have no such insights; the whole approach of reducible configurations, refined several times, is mathematically interesting. But checking a large number of cases by computer does not tell us whether there is some underlying simple reason why all those cases would work.

You might have guessed by now that I will not show you a proof of the four color theorem in this chapter. We will, however, look at the arguments involved in proving two weaker bounds.

24.3 Greedy coloring

The greedy coloring algorithm goes through the vertices and gives each one a color not already used on its neighbors. This algorithm is guaranteed to find a coloring of the graph, but the number of colors used depends on the order in which we color the vertices. The best thing we can say about the greedy algorithm in general is that if a graph G has maximum degree $\Delta(G)$, it will never use more than $\Delta(G) + 1$ colors.

Question: Does give us any kind of universal upper bound for planar graphs?

Answer: No: the maximum degree of a planar graph can be arbitrarily high. For example, the star graph S_n is a tree with $n-1$ leaf vertices adjacent to one central vertex; it is planar (like all trees) but has maximum degree $n-1$.

For planar graphs, it is possible to use some of the theory we've already developed to order the vertices more intelligently. The beginning is a bound on the minimum degree of a planar graph: though planar graphs can have vertices of very large degree, this cannot be true of every vertex.

Lemma 24.2. *Every planar graph G has minimum degree $\delta(G) \leq 5$.*

Proof. This is true for every planar graph with at most 6 vertices because at that point, you can't have any degrees bigger than 5.

For planar graphs with $n \geq 3$ vertices and m edges, Theorem 22.2 tells us that $m \leq 3n - 6$. However, if every vertex had degree 6 or more, then we would have $m \geq \frac{1}{2}(6n) = 3n$ by the handshake lemma (Lemma 4.1), and we cannot have $3n \leq 3n - 6$. Therefore not all vertices have degree 6 or more: there must be a vertex with degree 5 or less. \square

An alternate way to phrase the proof would be to look at the average degree of a vertex, rather than the total number of edges.

Question: What is the average degree of a graph with n vertices and $m \leq 3n - 6$ edges, and how does this help us?

Answer: It is given by $\frac{2m}{n} \leq \frac{2(3n-6)}{n}$, which simplifies to $6 - \frac{12}{n}$. Since $6 - \frac{12}{n} < 6$, the average degree is always less than 6. Not all vertices can be above average, so there must be a vertex of degree less than 6.

It is convenient to have a vertex of small degree, because we can leave it to be colored last: even if the worst should happen and all its neighbors have different colors, we will still be able to pick a color for it. For example, if we have 6 colors available, and we leave a vertex of degree 5 until the end, it will always be possible to give it a color.

Question: If $\delta(G) \leq 5$, is that enough to know that $\chi(G) \leq 6$: that G is 6-colorable?

Answer: No! For example, we could start with K_{100} and add a new vertex of degree 5 (or degree 0). That low-degree vertex can be left until the end, but we'll still need 100 colors to color K_{100} .

In the case of a planar graph, however, we know more. If we remove a vertex of minimum degree, what we're left with is a smaller planar graph, and Lemma 24.2 also applies to that smaller planar graph. That graph, too, has a vertex of degree 5 or less, which is enough to give us a proof by induction.

Lemma 24.3. *Every n -vertex planar graph G has a vertex ordering x_1, x_2, \dots, x_n in which each vertex is adjacent to at most 5 of the vertices that come before it.*

Proof. We will prove this by induction on n . When $n \leq 6$, any vertex ordering will do.

Assume that the lemma is true for all $(n - 1)$ -vertex planar graphs, and let G be an n -vertex planar graph. By Lemma 24.2, $\delta(G) \leq 5$; let x be a vertex with $\deg_G(x) \leq 5$.

Apply the induction hypothesis to find a vertex ordering x_1, x_2, \dots, x_{n-1} of $G - x$. We can extend it to a vertex ordering of G by setting $x_n = x$. For all $i \leq n - 1$, x_i has fewer than 5 neighbors among $\{x_1, x_2, \dots, x_{i-1}\}$ by the induction hypothesis. Meanwhile, x_n has fewer than 5 neighbors among $\{x_1, x_2, \dots, x_{n-1}\}$ because it has fewer than 5 neighbors total.

By induction, we can find such a vertex ordering for all n . □

Using this lemma, we can prove our first bound on the chromatic number of arbitrary planar graphs!

Theorem 24.4. *If G is a planar graph, then $\chi(G) \leq 6$.*

Proof. Let x_1, x_2, \dots, x_n be the vertex ordering given by Lemma 24.3, and let C be a set of 6 colors. For $i = 1, 2, \dots, n$ in that order, give x_i an arbitrary color in C not already used on its neighbors.

Why is this possible? Because x_i only has 5 neighbors in the set $\{x_1, x_2, \dots, x_{i-1}\}$, and these are the only vertices already given a color when we get to x_i . So at most 5 of the 6 colors in C have been used on the neighbors of x_i , and there is still a color left to choose.

When we're done, the coloring we get is proper, because we never give a vertex a color used on an adjacent vertex. For every edge $x_i x_j$, where $i < j$, we already knew the color of x_i when we got to x_j , and we made sure that x_j would be given a different color. \square

24.4 Five colors

Suppose we want to go one step further, and prove that $\chi(G) \leq 5$ for all planar graphs G .

Question: What would go wrong if we tried the methods in the previous section to prove this?

Answer: We would encounter planar graphs with minimum degree 5. Here, we cannot leave any vertex for last and forget about it; its neighbors might end up with 5 different colors, and we would have no color left to use on the last vertex.

Question: Are there, in fact, any planar graphs with minimum degree 5?

Answer: Yes, and we've seen them already: the skeleton graph of the icosahedron is one example.

We can still try to find an recursive algorithm for 5-coloring planar graphs: given a planar graph G and a vertex x of minimum degree, we 5-color $G - x$ and then complete the result to a coloring of G . However, if $\deg_G(x) = 5$, we will need to do one of two things:

1. Before we color $G - x$, modify it somehow to ensure that we'll be able to put back x and give it a color.
2. After we color $G - x$, modify the coloring somehow to ensure that we'll be able to put back x and give it a color.

Historically, the first proofs of Theorem 24.5 used the second approach, modifying the coloring using subgraphs called Kempe chains; this idea is similar to the recoloring strategy we used in the proof of Vizing's theorem (Theorem 20.6). Kempe chains are named after Alfred Kempe, who used the idea in 1879 in a proposed proof of the four color theorem. In 1890, Percy Heawood found a subtle error in Kempe's proof, but the argument via Kempe chains was still the first argument showing that five colors are always enough to color any map.

We will instead take the first approach, which is a shorter and more direct argument. This proof is due to Paul Kainen [15] and is a much later invention.

Theorem 24.5. *If G is a planar graph, then $\chi(G) \leq 5$.*

Proof. We induct on n , the number of vertices of G . (This will be a strong induction, because sometimes we'll need to go back from n to $n - 2$, not just $n - 1$.) If $n \leq 5$, then of course G has a coloring with at most 5 colors; this gives us our base cases.

Now assume that all planar graphs with fewer than n vertices have colorings with at most 5 colors, and let G be an n -vertex planar graph. By Lemma 24.2, G has a vertex x of degree at most 5.

If in fact $\deg(x) \leq 4$, then we have lucked out. By induction, $G - x$ has a coloring with at most 5 colors. In that coloring, the neighbors of G use at most 4 of the colors, because there are at most 4 of them, so we can give x a color none of them use. The result is a coloring of G .

If $\deg(x) = 5$, let y_1, y_2, y_3, y_4, y_5 be the neighbors of x . The first thing we observe is that it's not possible for all of the edges between y_1, \dots, y_5 to be present in G .

Question: Why not?

Answer: Then they'd form a copy of K_5 (a copy of K_6 , in fact, when taken together with x), which we know is not planar; but G is a planar graph.

So let y_i and y_j be two of x 's neighbors that are not adjacent in G . (We'll see why this matters soon!) Construct a graph H in the following steps:

1. Delete all edges from x except edges xy_i and xy_j .
2. Contract edges xy_i and xy_j , obtaining a new vertex z .

Deleting and contracting edges preserves planarity, so H is a planar graph on $n - 2$ vertices. By our induction hypothesis, H has a 5-coloring. Extend this 5-coloring to $G - x$ by giving y_i and y_j the color of z .

Question: Why is this a proper coloring?

Answer: Every neighbor of y_i or y_j in $G - x$ was a neighbor of z in H , so it had a different color from the color used on z . Therefore in $G - x$, it has a different color from the color used on y_i and y_j . This is all that needs to be checked; for all other edges, both endpoints have the same color in $G - x$ and in H .

Question: What would have gone wrong if y_i and y_j were adjacent?

Answer: Then both endpoints of the edge $y_i y_j$ would have the same color, so the coloring of $G - x$ would not be proper.




Figure 24.3: Using a Hamilton cycle in the borders of a map to 4-color it

Now we have a 5-coloring of $G - x$ in which only 4 different colors are used on the neighbors of x , because two of them (y_i and y_j) have the same color. Once again, this lets us give x a color not used on any of its neighbors and obtain a coloring of G .

This proves the induction step, and so the theorem is true for planar graphs with any number of vertices. \square

24.5 Hamilton cycles

In 1880, Peter Tait proposed several solutions of his own to the map coloring problem. At the time, Kempe's 1879 proof was generally accepted; Tait merely felt that the proof was too complicated, and wanted a simpler one. None of Tait's solutions worked out in the sense of giving an alternative proof, but several of them were successful in giving alternate avenues of attack on the problem.

Earlier, I mentioned that the graph of adjacencies of a map is the dual of a plane embedding whose edges are exactly the borders drawn in the map. However, we did not really consider this plane embedding as an object of study in its own right. Tait was the first to do so, and was able to get remarkably close to a proof of the four color theorem in this way.

Tait began by finding a Hamilton cycle in this plane embedding; for an example, consider the Hamilton cycle in Figure 24.3a. The Hamilton cycle divides the plane embedding into two parts, an inside and an outside. Tait's strategy was to color those two parts separately (see Figure 24.3b and Figure 24.3c) then combine the colorings, as shown in Figure 24.3d.

Why would this help? Well, Tait noticed that the inside and the outside of the cycle can each be colored with just two colors: four total! In fact, the graph of adjacencies on either side of the Hamilton cycle is a tree, and as we know (Proposition 13.3), all trees are bipartite, or in other words 2-colorable.

To see this, let H (for "half") be the graph of adjacencies on just one side of the Hamilton cycle. Each edge in H exists due to a border in the plane embedding where we drew the Hamilton cycle; both endpoints of that border lie on the cycle, because it's a Hamilton cycle. If a border is drawn between two points on a closed loop, it cuts that loop in half, separating the two parts of the side it's drawn on. Back in H , that makes the edge we were looking at a bridge. A connected graph in which every edge is a bridge is a tree.

Question: Why is this not a proof of the four color theorem?

Answer: The missing detail is that we don't yet have a reason to believe that the Hamilton cycle we are relying on exists!

In fact, it's pretty easy to draw a map in which we can't draw a Hamilton cycle along the borders. (Switzerland is an example: here, the borders aren't even connected!) Tait was aware of this, but still had hope to do it in all the worst cases, which would be enough to prove the four color theorem.

What are the worst cases? The phrase “worst case” is one you should generally watch out for in your proofs: if you're not careful, it can lead to unjustified assumptions. To properly use it, we should first explain how every other case is simpler than some “worst” case.

That's exactly what we can do here. Suppose we are trying to color an arbitrary planar graph G . If there is any edge xy such that adding xy to G produces another planar graph $G + xy$, then we might as well add it! Every coloring of $G + xy$ is also a coloring of G , with the extra restriction that x and y cannot be given the same color. And if we think we have a rule for coloring all planar graphs, that rule should be able to handle $G + xy$ just as well as G .

This means that it's enough to consider only maximal planar graphs, which cannot gain any new edge and still be planar. By Proposition 22.4, these are exactly the graphs whose plane embeddings are always triangulations.

Question: What do we know about the dual of a triangulation?

Answer: The condition that each face has length 3 turns into a condition that each vertex has degree 3. They are 3-regular graphs!

(In fact, not all 3-regular planar graphs are the duals of triangulations; the graph must also be 3-*connected*, but we won't learn what that means until Chapter 26.)

Tait's conjecture was that the dual graph of every triangulation is Hamiltonian. If this were true, it would imply the four color theorem, by the strategy we pursued in Figure 24.3a. This seemed to be a viable strategy for a long time; it was not until 1946 that Tutte found a counterexample [22]. That counterexample is shown in Figure 24.4 and is known as the *Tutte graph*. (Here, the planar graph which cannot be colored using Tait's strategy is the dual of the Tutte graph; that is, we wish to color the faces in Figure 24.4. By the four color theorem, of course, this is still possible—just not via finding a Hamilton cycle.)

24.6 Coloring empires

In 1890, along with finding the flaw in Kempe's attempt at the four color theorem, Heawood asked [12]: what about maps like the map of Switzerland? That is, what about maps where a single country can have multiple disconnected parts, which must be given the same color?

If we don't put some kind of restriction on this ability, then there is no limit to the number of colors we might need.




Figure 24.4: The Tutte graph

Question: Why not?

Answer: For example, give each country a tiny exclave in the middle of each other country's territory. Now this prevents those two countries from sharing a color, so every country will need a different color.

Heawood's approach was to express the problem in terms of the maximum number of parts a single country can have. Unfortunately, he did not come up with a clever name for such divided countries. This omission was corrected by Martin Gardner when writing about Heawood's work much later [8]; Gardner proposed the term *M-pire* for an empire whose territory consists of at most M disconnected pieces.

Let's begin by looking at the case $M = 2$. (This more or less describes Switzerland; though some Swiss cantons have more than 2 pieces, this does not contribute any additional adjacencies.)

Proposition 24.6. *A map of 2-pires can always be colored using at most 12 colors so that two 2-pires which share a border (in at least one of their territories) receive different colors.*

Proof. We can model a map of 2-pires in two ways:

1. With a graph (call it H) in which vertices are connected regions, and an edge exists whenever two regions share a border. Let there be n regions; then (as we've already seen) this graph is planar and has at most $3n - 6$ edges.
2. With a graph (call it G) in which vertices are the empires, and an edge exists between two vertices if some territory belonging to one empire borders some territory belonging to the other empire. This is the graph we actually want to color.

In G , there are also at most $3n - 6$ edges! That's because every edge in G must come from an edge in H (a border shared is a border shared in either representation) but every edge in H corresponds to at most one edge in G (a shared border only belongs to two empires).

Question: Is it possible for G to have fewer edges than H ?

Answer: Yes: whenever one empire touches two of another empire's regions, that is represented by two edges in H , but only one edge in G .




Figure 24.5: Heawood's map of 12 pairwise adjacent empires

Question: What can we say about the number of vertices in G ?

Answer: It could be as high as n (if every 2-pire is actually an ordinary country), but cannot go below $n/2$: with at most 2 regions per empire, it takes at least $n/2$ empires to reach n regions.

The average degree in G is $\frac{2|E(G)|}{|V(G)|}$ by the handshake lemma, or at most $\frac{3n-6}{n/2} = 12 - \frac{24}{n}$. In particular, it is less than 12, so $\delta(G) < 12$.

What's more, if we delete a vertex x from G , then the remaining graph $G - x$ is also the graph of adjacencies of some map of 2-pires: just erase empire x from the map! This means that $\delta(G - x) < 12$, by the same argument, and in general, every subgraph of G will have a vertex of degree less than 12.

We can now color G by the same greedy argument as we used to prove Theorem 24.4, using 12 colors instead of 6. Leaving a low-degree vertex until the end guarantees that one of 12 colors will be available for it, since its neighbors have at most 11 different colors. \square

Theorem 24.4 was considerably more pessimistic than the truth: it says that $\chi(G) \leq 6$ for all planar graphs G , but actually, it is also true that $\chi(G) \leq 4$. Does Proposition 24.6 have a similar flaw?

No: it turns out that it is the best possible bound! Figure 24.5 shows a map drawn by Heawood. Here, there are 12 empires (labeled 1 through 12), and any two of them share a border! For this map, it is impossible to use fewer than 12 colors, because all empires must have different colors; therefore there cannot be a general result better than Proposition 24.6.

Heawood gave an argument for the general case of M -pires, as well.

Question: What upper bound does the same argument give for a map of M -pires?

Answer: An upper bound of $6M$ colors: with n regions, there are still at most $3n - 6$ edges (less than $3n$) but the number of M -pires is at most n/M , so the average degree is less than $\frac{2(3n)}{n/M} = 6M$.

However, he was unable to find a map to show that this was the best upper bound when $M \geq 3$, and left this as a conjecture. This conjecture was resolved almost a century later, when Brad Jackson and Gerhard Ringel showed how to construct such maps for all $M \geq 2$ [14].


24.7 Practice problems

1. a) Find a 4-coloring of the faces of an icosahedron.
b) Prove that 3 faces are not enough to color the faces of an icosahedron.
2. Even though Switzerland's graph of adjacencies is not planar, it is still 4-colorable! Find such a 4-coloring. (You may prefer to refer to the diagram in Figure 1.2 all the way back in the first chapter instead of a map of Switzerland.)
3. The mathematician August Ferdinand Möbius is perhaps best known for his mathematical study of the Möbius strip: the surface formed by taking a strip of paper and joining the two ends with a half-twist, so that the two sides of the paper are turned into one side. In regards to coloring maps, his contribution was to show that 5 connected regions on a map cannot all share borders with each other; an early form of proving that K_5 is not planar.

Ironically, when drawing a map on a Möbius strip rather than in the plane, it is possible to have 5 regions all border each other. How can this be done?


As a follow-up, see if you can do it for 6 regions.

4. The Tutte graph shown in Figure 24.4 is made up of three subgraphs isomorphic to the one shown below, plus a central vertex. The three vertices labeled x , y , and z in the diagram are used to connect the three subgraphs together: each x -vertex is joined to a y -vertex in a different subgraph to connect the three cyclically, and the three z -vertices are all joined to the center vertex of the Tutte graph.



- a) Show (by casework) that the graph above has no $x - y$ Hamilton path. (It does have $x - z$ and $y - z$ Hamilton paths, which you may feel free to find, though they won't be useful for the next step.)
- b) Use part (a) to show that the Tutte graph has no Hamilton cycle.

5. Here are several coloring problems for the union of two graphs. The graphs may share vertices and even edges; their union is a graph containing every vertex and every edge present in at least one of the graphs.
 - a) The union of two planar graphs is 12-colorable. Explain why this is a special case of Proposition 24.6.
 - b) By imitating the proof of Proposition 24.6, prove that the union of two trees is always 4-colorable.
 - c) Prove that the union of two bipartite graphs is always 4-colorable.
6. The faraway and fictional continent of Kvadrat is divided into many countries, all in the shape of 1×1 squares. The squares are not necessarily aligned to a grid, and might have unclaimed land between them (which does not need to be given a color). One possible map is shown below:



Can all possible maps of Kvadrat be colored using only three colors, or is there an example of such a map which requires four colors?

7. The faraway and fictional country of Heibai is divided into cantons, just like Switzerland, but its map has a curious property: at every point where several borders meet, the total number of borders that converge there is even. (The United States has just one point of this type: the point where Arizona, Colorado, New Mexico, and Utah come together.)
 Prove that it's possible to color the cantons of Heibai with just two colors so that two cantons that share a border are always colored differently. (Two cantons that only share a point may have the same color.)
8. Let G be a planar graph that has been partially 5-colored: there is a set $W \subseteq V(G)$ such that every vertex $x \in W$ has already been given a color $f(x) \in \{1, 2, 3, 4, 5\}$, which cannot be changed.

Prove that if the distance between any two vertices in W is at least 4, then this partial coloring can be completed to a 5-coloring of G . (This result was proved by Michael Albertson in 1998 [1].)

Bibliography

- [1] Michael Albertson. “You can’t paint yourself into a corner”. In: *Journal of Combinatorial Theory, Series B* 73 (2 1998), pp. 189–194. DOI: doi.org/10.1006/jctb.1998.1827.
- [2] Kenneth Appel and Wolfgang Haken. “Every planar map is four colorable. Part I: Discharging”. In: *Illinois Journal of Mathematics* 21 (3 1977), pp. 429–490. DOI: [10.1215/ijm/1256049011](https://doi.org/10.1215/ijm/1256049011).
- [3] Kenneth Appel and Wolfgang Haken. “Every planar map is four colorable. Part II: Reducibility”. In: *Illinois Journal of Mathematics* 21 (3 1977), pp. 491–567. DOI: [10.1215/ijm/1256049012](https://doi.org/10.1215/ijm/1256049012).
- [4] K. Coolsaet, S. D’hont, and J. Goedgebeur. *House of Graphs 2.0: A database of interesting graphs and more*. 2023. DOI: [10.1016/j.dam.2022.10.013](https://doi.org/10.1016/j.dam.2022.10.013). URL: <https://houseofgraphs.org>.
- [5] David Eppstein. *Twenty-one Proofs of Euler’s Formula: $V - E + F = 2$* . URL: <https://ics.uci.edu/~eppstein/junkyard/euler/> (visited on 10/17/2025).
- [6] Leonhard Euler. “Elementa doctrinae solidorum”. In: *Novi Commentarii academiae scientiarum Petropolitanae* 4 (1758), pp. 109–140. URL: <https://scholarlycommons.pacific.edu/euler-works/230/>.
- [7] P.J. Federico. *Descartes on Polyhedra: A Study of the “De Solidorum Elementis”*. Sources in the History of Mathematics and Physical Sciences. Springer, 1982. ISBN: 978-0-387-90760-4.
- [8] Martin Gardner. “M-Pire Maps”. In: New York, NY: Springer New York, 1997, pp. 85–100. ISBN: 978-0-387-30389-5.
- [9] Martin Gardner. “The Five Platonic Solids”. In: New York, NY: Cambridge University Press, 2008, pp. 1–10. ISBN: 978-0-521-75610-5.
- [10] Georges Gonthier. “Formal Proof—The Four-Color Theorem”. In: *Notices of the American Mathematical Society* 55 (11 2008), pp. 1382–1393.
- [11] Richard Guy. “The strong law of small numbers”. In: *The American Mathematical Monthly* 95 (8 1988), pp. 697–712. DOI: [10.1080/00029890.1988.11972074](https://doi.org/10.1080/00029890.1988.11972074).
- [12] P.J. Heawood. “Map colour problem”. In: *Quarterly Journal of Pure and Applied Mathematics* 24 (1890), pp. 332–338.
- [13] Hud Hudson. “Four Colors Do Not Suffice”. In: *The American Mathematical Monthly* 110 (5 2003), pp. 417–423. DOI: [10.1080/00029890.2003.11919979](https://doi.org/10.1080/00029890.2003.11919979).
- [14] Brad Jackson and Gerhard Ringel. “Solution of Heawood’s empire problem in the plane”. In: *Journal für die Reine und Angewandte Mathematik* 347 (1984), pp. 146–153. DOI: [10.1515/crll.1984.347.146](https://doi.org/10.1515/crll.1984.347.146).
- [15] Paul C. Kainen. “A generalization of the 5-color theorem”. In: *Proceedings of the American Mathematical Society* 45 (3 1974), pp. 450–453. DOI: [10.2307/2039977](https://doi.org/10.2307/2039977).

- [16] David Kullman. “The Utilities Problem”. In: *Mathematics Magazine* 52 (5 1979), pp. 299–302. DOI: [10.1080/0025570X.1979.11976807](https://doi.org/10.1080/0025570X.1979.11976807).
- [17] Casimir Kuratowski. “Sur le problème des courbes gauches en topologie”. In: *Fundamenta Mathematicae* 15 (1930), pp. 271–283. DOI: [10.4064/fm-15-1-271-283](https://doi.org/10.4064/fm-15-1-271-283).
- [18] Mikhail Lavrov. *Sequence A380996 in the On-Line Encyclopedia of Integer Sequences*. URL: <https://oeis.org/A380996> (visited on 10/17/2025).
- [19] Neil Robertson et al. “The Four-Colour Theorem”. In: *Journal of Combinatorial Theory, Series B* 70 (1 1997), pp. 2–44. DOI: [10.1006/jctb.1997.1750](https://doi.org/10.1006/jctb.1997.1750).
- [20] SimpleMaps.com. *Free GIS Maps of Switzerland*. Licenced under [CC BY 4.0](https://creativecommons.org/licenses/by/4.0/). 2010–2025. URL: <https://simplemaps.com/gis/country/ch> (visited on 10/26/2025).
- [21] N.J.A. Sloane. *Sequence A000109 in the On-Line Encyclopedia of Integer Sequences*. URL: <https://oeis.org/A000109> (visited on 10/17/2025).
- [22] William Thomas Tutte. “On Hamiltonian circuits”. In: *Journal of the London Mathematical Society* 21 (2 1946), pp. 98–101. DOI: [10.1112/jlms/s1-21.2.98](https://doi.org/10.1112/jlms/s1-21.2.98).
- [23] W.T. Tutte. “Matroids and Graphs”. In: *Transactions of the American Mathematical Society* 90 (3 1959), pp. 527–552. DOI: [10.2307/1993185](https://doi.org/10.2307/1993185).
- [24] Klaus Wagner. “Über eine Eigenschaft der ebenen Komplexe”. In: *Mathematische Annalen* 114 (1937), pp. 570–590. DOI: [10.1007/BF01594196](https://doi.org/10.1007/BF01594196).
- [25] Douglas B. West. *Introduction to Graph Theory*. Prentice Hall, 1996. ISBN: 9780132278287.
- [26] Robin J. Wilson. *Four colors suffice: how the map problem was solved*. Princeton, NJ: Princeton University Press, 2002. ISBN: 0691115338.