CS 70 Discrete Mathematics and Probability Theory
Summer 2019 Course Notes **Note 5.5**

1 Planarity, Euler's Formula, Coloring.

1.1 Planar Graphs

A graph is *planar* if it can be drawn on the plane without crossings. For example, the first four graphs shown below are planar. Notice that the first and second graphs are the same, but drawn differently. Even though the second drawing has crossings, the graph is still considered planar since it is possible to draw it without crossings.

The other three graphs are not planar. The first one of them is the infamous "three houses-three wells graph," also called $K_{3,3}$. The second is the complete graph with five nodes, or K_5. The third is the four-dimensional cube. We shall soon see how to prove that all three graphs are non-planar.

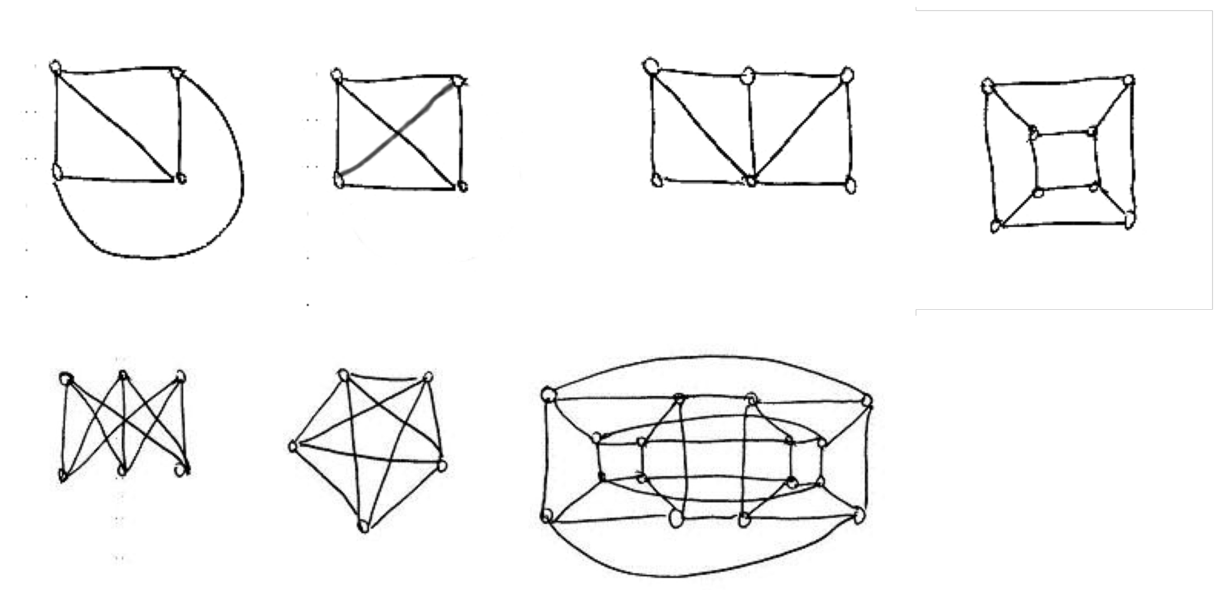

When a planar graph is drawn on the plane, one can distinguish, besides its vertices (their number will be denoted v here) and edges (their number is e), the *faces* of the graph (more precisely, of the drawing). The faces are the regions into which the graph subdivides the plane. One of them is infinite, and the others are finite. The number of faces is denoted f. For example, for the first graph shown $f = 4$, and for the fourth (the cube) $f = 6$.

There is a natural connection between faces and cycles, which we use to prove the following lemma:

Lemma 5.3. *Let G be a connected, planar graph. Then G is a tree if and only if G has a single face.*

Proof. We first prove the "if" direction by contraposition. Suppose that G is not a tree. Because G is connected but is not a tree, it must contain a cycle; consider the shortest cycle in G. The area on the "inside" of this cycle is not connected to the area on the "outside", and thus must be part of different faces. Hence, G must have more than one face.

Now we prove the "only if" direction, again by contraposition. Suppose that G has more than one face. Since there can only be a single infinite face, this must mean that there exists a finite face. If we trace the boundary of one such face, this will give us a cycle in G.[1] Thus, G has a cycle, and hence is not a tree. □

One basic and important fact about planar graphs is *Euler's formula*, $v + f = e + 2$ (check it for the graphs above). It has an interesting story. The ancient Greeks knew that this formula held for all polyhedra (check it for the cube, the tetrahedron, and the octahedron, for example), but could not prove it. How do you do induction on polyhedra? How do you apply the induction hypothesis? What is a polyhedron minus a vertex, or an edge? In the 18th century Euler realized that this is an instance of the inability to prove a theorem by induction *because it is too weak*, something that we saw time and again when we were studying induction. To prove the theorem, one has to generalize polyhedra. And the right generalization is *planar graphs*.

Sanity check! Can you see why planar graphs generalize polyhedra? Why are all polyhedra (without "holes") planar graphs?

Theorem 5.3. *(Euler's formula) For every connected planar graph,* $v + f = e + 2$

Proof. We prove this by induction on f.

Base Case ($f = 1$): By Lemma 5.3, our graph is a tree. We know that a tree has one fewer edge than vertices, so $e = v - 1$. Hence, $e + 2 = v + 1 = v + f$, as desired.

Inductive Step: Suppose that this holds for all planar graphs with k faces. Let $G = (V, E)$ be any planar graph with $k + 1$ faces. Suppose now we remove an edge that is on the boundary between two faces, resulting in a graph $G' = (V, E')$. This cannot disconnect the graph, as that would require that the infinite face was originally on both sides of the removed edge, contradicting that it was a boundary between two faces. Additionally, removing that edge will combine the two faces it was a boundary between. Thus, G' will be a connected graph with k faces, so by the inductive hypothesis $|V| + k = |E'| + 2$. But $|E'| = |E| - 1$, so adding one to both sides, we get $|V| + (k + 1) = |E| + 2$, showing that Euler's formula holds for G.

□

Sanity check! What happens when the graph is not connected? How does the number of connected components enter the formula?

Take a planar graph with f faces, and consider one face. It has a number of *sides*, that is, edges that bound it clockwise. Note that an edge may be counted twice, if it has the same face on both sides, as it happens for example in a tree (such edges are called bridges). Denote by s_i the number of sides of face i. Now, if we add the s_i's we are going to get $2e$, because each edge is counted twice, once for the face on its right and once for the face on its left (they may coincide if the edge is a bridge). We conclude that, in any planar graph,

$$\sum_{i=1}^{f} s_i = 2e. \qquad (1)$$

Now notice that, since we don't allow parallel edges between the same two nodes, and if we assume that there are at least two edges (so there are at least three vertices), every face has at least three sides, or $s_i \geq 3$

[1] There is a slight caveat to this, which is that if there is a path sticking out into the face (called a *bridge*), we will try to reuse all those edges twice. However, we can fix this by simply ignoring all such edges when tracing the boundary.

for all i. It follows that $3f \leq 2e$. Solving for f and plugging into Euler's formula we get

$$e \leq 3v - 6.$$

This is an important fact. First it tells us that planar graphs are *sparse*, they cannot have too many edges. A 1,000-vertex connected graph can have anywhere between a thousand and half a million edges. This inequality tells us that for planar graphs the range is very small, between 999 and $2,994$.

It also tells us that K_5 is not planar: Just notice that it has five vertices and ten edges.

$K_{3,3}$ has $v = 6, e = 9$ so it passes the planarity test with flying colors. We must think a little harder to show that $K_{3,3}$ is non-planar. Notice that, if we had drawn it on the plane, there would be no triangles. Because a triangle means that two wells or two houses are connected together, which is false. So, Equation (1) now gives us $4f \leq 2e$, and solving for f and plugging into Euler's formula, $e \leq 2v - 4$, which shows that $K_{3,3}$ is non-planar.

So, we have established that K_5 and $K_{3,3}$ are both non-planar. There is something deeper going on: In some sense, these are *the only non-planar graphs*. This is made precise in the following famous result, due to the Polish mathematician Kuratowski (this is what "K" stands for).

Theorem 5.4. *A graph is non-planar if and only if it contains K_5 or $K_{3,3}$.*

"Contains" here means that one can identify nodes in the graph (five in the case of K_5, six in the case of $K_{3,3}$) which are connected as the corresponding graph through paths (possibly single edges), and such that no two of these paths share no vertex. For example, the 4-cube shown below is non-planar, because it contains $K_{3,3}$, as shown.

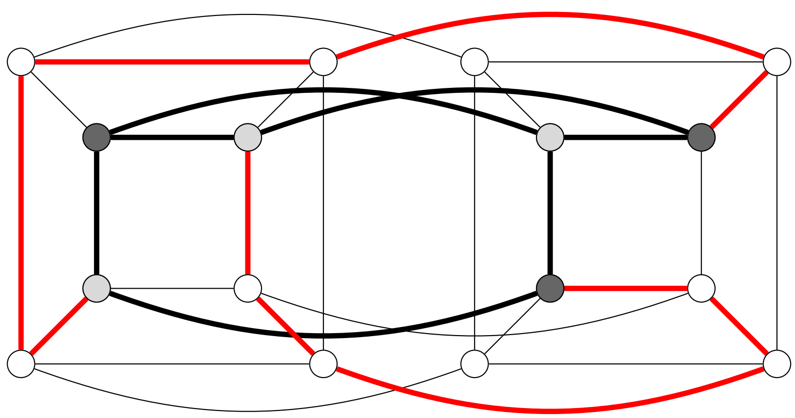

Can you find K_5 in the same graph?

One direction of Kuratowski's theorem is obvious: If a graph contains one of these two non-planar graphs, then of course it is itself non-planar. The other direction, namely that in the absence of these graphs we can draw any graph on the plane, is difficult. For a short proof you may want to type "proof of Kuratowski's theorem" in your favorite search engine.

1.2 Duality and Coloring

There is an interesting *duality* between planar graphs. For example, the Greeks knew that the octahedron and the cube are "dual" to each other, in the sense that the faces of one can be put in correspondence with the vertices of the other (think about it). The tetrahedron is self-dual. And the dodecahedron and the icosahedron (look for images in the web if you don't know these pretty things) are also dual to one another.

What does this mean? Take a planar graph G, and assume it has no bridges and no degree-two nodes. Draw a new graph G^*: Start by placing a node on each face of G. Then draw an edge between two faces if they touch at an edge — draw the new edge so that it crosses that edge. The result is G^*, also a planar graph. Notice now that, if you construct the dual of G^*, it is the original graph: $(G^*)^* = G$.

Duality is a convenient consideration when thinking about planar graphs. Also, it tells us that "coloring a political map so that no two countries who share a border have the same color" is the same problem as "coloring the vertices of a planar graph (the dual of the political map) so that no two adjacent vertices have the same color."

A famous theorem states that four colors are always enough! (Search for "four color theorem".)

We shall prove something weaker:

Theorem 5.5. *Every planar graph can be colored with five colors.*

Before proving this theorem, we will first prove an even weaker version, which will give us ideas about how to proceed:

Theorem 5.6. *Every planar graph can be colored with six colors.*

Proof. We prove this statement by induction on $|V|$.

Base Case ($|V| = 1$): In this case, our graph has a single vertex, and so can be colored using only 1 color.

Inductive Step: Suppose that any graph on k vertices can be 6-colored, and let $G = (V,E)$ be a graph on $k+1$ vertices. We've previously shown that for any planar graph, $e \leq 3v - 6$, so in particular $|E| \leq 3|V| - 6$. Since the sum of the degrees of every vertex is $2|E|$, we have that this total degree is at most $6|V| - 12$. Hence, the average degree of a vertex in G is $\frac{6|V|-12}{|V|} < 6$. Since it's impossible for every vertex to have above-average degree, we must have that there exists a degree 5 or less vertex in G.

Let v be such a vertex. If we remove v, we're left with a (planar) graph on $|V| - 1$ vertices, and so can color it using the inductive hypothesis. If we now add v back in, we know that it will have at most 5 neighbors. Even if all these neighbors are different colors, there is still a 6th color remaining for us to color v with. If we assign v this color and leave everything else the same, we have a valid coloring for G, completing the inductive step. □

The proof that five colors suffices follows much the same argument as above. The main additional challenge is that, when adding back the vertex we removed, it is possible that all 5 of its neighbors have different colors, leaving none left to color the readded vertex with. In order to get around this, we will need to have some way of slightly modifying the coloring given by the inductive hypothesis.

Following this line of reasoning, we consider the subset of vertices colored one of two colors, say 1 and 2, and compute the connected components of the result. We note that we can flip the two colors within such a connected component — any edge that is not in the connected component is clearly fine, and any edge in the connected component has both endpoints switched which is also fine.

Proof of Theorem 5.5. We again proceed by induction of $|V|$.

Base Case ($|V| = 1$): In this case, our graph has a single vertex, and so can be colored using only 1 color.

Inductive Step: Suppose that any planar graph on k vertices can be 5-colored, and let $G = (V,E)$ be a planar graph on $k+1$ vertices. As before, we find some vertex v in G such that $\deg(v) \leq 5$, remove it, and use the inductive hypothesis to color the result. If the degree of v is 4 or less, or if two of v's neighbors are given the same color, we immediately have a color available to give to v, and so are done.

The only case now left to consider is if v has 5 neighbors, all of which are given different colors by the inductive hypothesis. Fix some planar drawing of G, and label those colors c_1 through c_5 as depicted below:

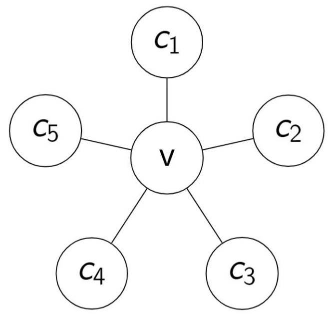

We first consider the graph using only vertices with colors c_5 and c_3. Consider the connected component of v's c_5-colored neighbor in this graph. If it does not contain v's c_3-colored neighbor, we can flip all of the colors in that component, leaving c_5 available for v. Otherwise, we must have a path from the c_5 neighbor to the c_3 one consisting only of vertices colored c_3 or c_5.

If this fails to give us a color for v, we next try the same trick with the c_2- and c_4-colored neighbors of v. As before, the only way this can fail to give us a color for v is if there is a path connecting those two neighbors consisting only of vertices colored c_2 or c_4. Thus, the only case we have left to worry about is where we have a c_3/c_5 colored path and a c_2/c_4 colored one, as depicted below:

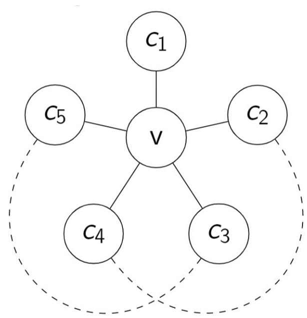

We can see that these two paths must intersect at some point. Because the drawing is planar, they must intersect at a vertex. But this vertex simultaneously must be colored c_3/c_5 and c_2/c_4, which is impossible. Thus, this case cannot actually happen, so the cases where we were able to find a color for v are exhaustive.

□

2 Hypercubes

We have discussed how complete graphs are a class of graphs whose vertices are particularly "well-connected." However, to achieve this strong connectivity, a large number of edges is required, which in many applications of graph theory is infeasible. Consider the example of the *Connection Machine*, which was a massively parallel computer by the company Thinking Machines in the 1980s. The idea of the Connection Machine was to have a million processors working in parallel, all connected via a communications network. If you were to connect each pair of such processors with a direct wire to allow them to communicate (i.e., if you used a complete graph to model your communications network), this would require 10^{12} wires! What the builders of the Connection Machine thus decided was to instead use a 20-dimensional *hypercube* to model their network, which still allowed a strong level of connectivity, while limiting the number of neighbors of each processor in the network to 20. This section is devoted to studying this particularly useful class of graphs, known as hypercubes.

The vertex set of the n-dimensional hypercube $G = (V, E)$ is given by $V = \{0,1\}^n$, where recall $\{0,1\}^n$ denotes the set of all n-bit strings. In other words, each vertex is labeled by a unique n-bit string, such as $00110\cdots0100$. The edge set E is defined as follows: Two vertices x and y are connected by edge $\{x,y\}$ if and only if x and y differ in exactly one bit position. For example, $x = 0000$ and $y = 1000$ are neighbors, but $x = 0000$ and $y = 0011$ are not. More formally, $x = x_1x_2\dots x_n$ and $y = y_1y_2\dots y_n$ are neighbors if and only if there is an $i \in \{1,\dots,n\}$ such that $x_j = y_j$ for all $j \neq i$, and $x_i \neq y_i$. To help you visualize the hypercube, we depict the 1-, 2-, and 3-dimensional hypercubes below.

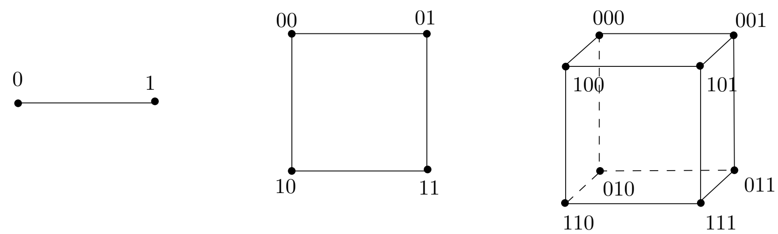

There is an alternative and useful way to define the n-dimensional hypercube via recursion, which we now discuss. Define the 0-subcube (respectively, 1-subcube) as the $(n-1)$-dimensional hypercube with vertices labeled by $0x$ for $x \in \{0,1\}^{n-1}$ (respectively, $1x$ for $x \in \{0,1\}^{n-1}$). Then, the n-dimensional hypercube is obtained by placing an edge between each pair of vertices $0x$ in the 0-subcube and $1x$ in the 1-subcube.

Sanity check! Where are the 0- and 1-subcubes in the 3-dimensional hypercube depicted above? Can you use these along with the recursive definition above to draw the 4-dimensional hypercube?

Exercise. Prove that the n-dimensional hypercube has 2^n vertices. Hint: Use the fact that each bit has two possible settings, 0 or 1.

We began this section by singing praises for the hypercube in terms of its connectivity properties; we now investigate these claims formally. Let us begin by giving two proofs of a simple property of the hypercube. Each proof relies on one of our two equivalent (namely, direct and recursive) definitions of the hypercube.

Lemma 5.4. *The total number of edges in an n-dimensional hypercube is $n2^{n-1}$.*

Proof 1. The degree of each vertex is n, since n bit positions can be flipped in any $x \in \{0,1\}^n$. Since each edge is counted twice, once from each endpoint, this yields a total of $n2^n/2 = n2^{n-1}$ edges. □

Proof 2. By the second definition of the hypercube, it follows that $E(n) = 2E(n-1) + 2^{n-1}$, and $E(1) = 1$, where $E(n)$ denotes the number of edges in the n-dimensional hypercube. A straightforward induction shows that $E(n) = n2^{n-1}$. □

Exercise. Using induction to show that in Proof 2 above, $E(n) = n2^{n-1}$.

Let us focus on the question of connectivity, and prove that the n-dimensional hypercube is well-connected in the following sense: To disconnect any subset $S \subseteq V$ of vertices from the rest of the graph, a large number of edges must be discarded. In particular, we shall see that the number of discarded edges must scale with $|S|$. In the theorem below, recall that $V - S = \{v \in V : v \notin S\}$ is the set of vertices that are not in S.

Theorem 5.7. *Let $S \subseteq V$ be such that $|S| \le |V - S|$ (i.e., that $|S| \le 2^{n-1}$), and let E_S denote the set of edges connecting S to $V - S$, i.e.,*

$$E_S := \{\{u,v\} \in E \mid u \in S \text{ and } v \in V - S\}.$$

Then, it holds that $|E_S| \ge |S|$.

Proof. We proceed by induction on n.

Base case ($n = 1$): The 1-dimensional hypercube graph has two vertices 0 and 1, and one edge $\{0,1\}$. We also have the assumption $|S| \le 2^{1-1} = 1$, so there are two possibilities. First, if $|S| = 0$, then the claim trivially holds. Otherwise, if $|S| = 1$, then $S = \{0\}$ and $V - S = \{1\}$, or vice versa. In either case we have $E_S = \{0,1\}$, so $|E_S| = 1 = |S|$.

Inductive step: We suppose the statement is true for $n = k$, and wish to prove the claim for $n = k+1$. Recall that we have the assumption $|S| \le 2^k$. Let S_0 (respectively, S_1) be the vertices from the 0-subcube (respectively, 1-subcube) in S. We have two cases to examine: Either S has a fairly equal intersection size with the 0- and 1-subcubes, or it does not.

1. **Case 1:** $|S_0| \le 2^{k-1}$ **and** $|S_1| \le 2^{k-1}$

   In this case, we can apply the induction hypothesis separately to the 0- and 1-subcubes. This says that restricted to the 0-subcube itself, there are at least $|S_0|$ edges between $|S_0|$ and its complement (in the 0-subcube), and similarly there are at least $|S_1|$ edges between $|S_1|$ and its complement (in the 1-subcube). Thus, the total number of edges between S and $V - S$ is at least $|S_0| + |S_1| = |S|$.

2. **Case 2:** $|S_0| > 2^{k-1}$

   In this case, S_0 is too large for the induction hypothesis to apply. However, note that since $|S| \le 2^k$, we have $|S_1| = |S| - |S_0| \le 2^{k-1}$, so we *can* apply the hypothesis to S_1. As in Case 1, this allows us to conclude that there are at least $|S_1|$ edges in the 1-subcube crossing between S and $V - S$.

   What about the 0-subcube? Here, we cannot apply the induction hypothesis directly, but there is a way to apply it after a little massaging. Consider the set $V_0 - S_0$, where V_0 is the set of vertices in the 0-subcube. Note that $|V_0| = 2^k$ and $|V_0 - S_0| = |V_0| - |S_0| = 2^k - |S_0| < 2^k - 2^{k-1} = 2^{k-1}$. Thus, we *can* apply the inductive hypothesis to the set $V_0 - S_0$. This yields that the number of edges between S_0 and $V_0 - S_0$ is at least $2^k - |S_0|$. Adding our totals for the 0-subcube and the 1-subcube so far, we conclude there are at least $2^k - |S_0| + |S_1|$ crossing edges between S and $V - S$. However, recall our goal was to show that the number of crossing edges is at least $|S|$; thus, we are still short of where we wish to be.

   But there are a still edges we have not accounted for — namely, those in E_S which cross between the 0- and 1-subcubes. Since there is an edge between every vertex of the form $0x$ and the corresponding vertex $1x$, we conclude there are at least $|S_0| - |S_1|$ edges in E_S that cross between the two subcubes. Thus, the total number of edges crossing is at least $2^k - |S_0| + |S_1| + |S_0| - |S_1| = 2^k \ge |S|$, as desired.

□