

6

Extremal combinatorics

6.1 Extremal graph theory

Real life is full of optimization problems, that is, we often have to find the best way to carry out a certain task. This often involves finding the graph that is best from a specific viewpoint. This is the topic of Extremal Graph Theory.

6.1.1 Bipartite graphs

In the previous chapter, we saw graphs whose vertices were colored so that there were no edges between vertices of the same color. If we have only one color, then such a coloring will exist only for empty graphs. So the first non-trivial case is when we have two colors. This case is very interesting on its own, which is why it has its own name.

Definition 6.1 *A graph G is called bipartite if its set of vertices has a proper 2-coloring.*

Bipartite graphs abound in everyday life. Imagine, for instance, a company that has some job openings and some candidates for them. A candidate may be qualified for more than one job, and a job may attract more than one candidate. The information relevant to this situation can be represented by a bipartite graph in which the red vertices are the jobs, the blue vertices are the candidates, and there is an edge between two vertices if the candidate associated with the blue vertex is qualified for the job associated to the red vertex. This definition assures that there will be no edges between vertices of the same color, so the obtained graph will be bipartite.

Bipartite graphs can be defined without resorting to colors, using the following fact.

Proposition 6.2 *A graph G is bipartite if and only if it does not contain a cycle of odd length.*

Proof: Let G be bipartite and let f be a proper 2-coloring of the vertices of G . If G had a cycle C of length $2k + 1$, then C would contain at least $k + 1$

vertices of the same color, so C would contain two adjacent vertices of the same color.

Conversely, assume that G has no odd cycles. Then start coloring the vertices of G at a vertex V , coloring it red, then coloring its neighbors blue, then coloring the neighbors of those red, and so on until all vertices of G are colored. This procedure will always lead to a proper 2-coloring, since vertices that are at even distance from V are red, and vertices that are at odd distance are blue. So if there were an edge between vertices A and B of the same color, then there would be an odd cycle in G . Indeed, take a shortest path p from V to A , and the shortest path p' from V to B . Note that A and B are at equal distances from V since they are neighbors and have the same color. If $p \cap p' = V$, then the union of path p , path p' , and edge AB forms an odd cycle. If not, let W be a common vertex of p and p' that is as far away from V as possible. Then A and B are at equal distance from W , and so the part of p that is between A and W , the part of p' that is between B and W , and AB together form an odd cycle. See Figure 6.1 for an illustration. \diamond




Figure 6.1

If there are no odd cycles, then our 2-coloring is proper.

Before we go any further, we should clarify what we mean by a *subgraph* of a graph, since this notion will be essential for the rest of this section. There are two different concepts as follows.

Definition 6.3 *Let G be a graph. We say that the graph H is a subgraph of G if the vertex set of H is a subset of the vertex set of G and the edge set of H is a subset of the edge set of G .*

Fair enough, you could say, but what else is there to say about defining subgraphs? The following definition—and the example after it—will hopefully answer that question.

Definition 6.4 *Let G be a graph. We say that the graph H is an induced subgraph of G if the vertex set $V(H)$ of H is a subset of the vertex set of G and if, when we take all edges of G that connect two vertices of $V(H)$, we get H itself.*

The following example should clarify the difference between these two notions.

Example 6.5 *Graph G in Figure 6.2 contains several subgraphs isomorphic to graph H . Indeed, any path of two edges is such a subgraph. However, G contains no induced subgraphs isomorphic to H since any 3-vertex induced subgraph of G is a copy of K_3 .*




Figure 6.2

The graph G contains a copy of H as a subgraph, but not as an induced subgraph.

No matter which definition of subgraph graphs we use, we see that, if G is bipartite, then so is any subgraph, or induced subgraph, of G . Conversely, if we keep adding edges to G while not adding new vertices, then eventually G will lose the bipartite property, since the complete graph on n vertices is not bipartite if $n \geq 3$. The question is *when* this will happen.

On one hand, trees are certainly bipartite graphs since they do not have any cycles, even or odd. So a bipartite graph on n vertices can certainly have $n - 1$ edges. On the other hand, since the complete graph K_n is not bipartite for $n \geq 3$, bipartite graphs on n vertices have fewer than $\binom{n}{2}$ edges. Where is the breaking point between $n - 1$ and $\binom{n}{2}$? The following result shows that it is closer to $\binom{n}{2}$ than to $n - 1$.

Theorem 6.6 *Let $n \geq 2$ and let G be a simple bipartite graph on n vertices. Then G has at most $a(n - a)$ edges, where $a = \lfloor n/2 \rfloor$.*

In other words, a bipartite graph on n vertices cannot have more than $n^2/4$ edges.

Proof: Assume G has a red vertices and $n - a$ blue vertices. Then the highest number of edges is achieved if all red vertices are adjacent to all blue vertices, which allows G to have $a(n - a)$ edges. By elementary calculus, the function $f(a) = a(n - a)$ has its maximum at $a = n/2$, and it increases when $0 < a < n/2$ and decreases when $a > n/2$. Because a has to be an integer, our claim follows. \diamond

A few remarks are in order. First, the graph described in the proof above, that is, the bipartite graph in which there are a red vertices, $n - a$ blue vertices, and *all red vertices* are adjacent to *all blue vertices*, is called a *complete bipartite graph*, and is denoted by $K_{a,n-a}$. So, for instance, the graph shown in Figure 5.53 is $K_{2,3}$.

Second, we have proved a little bit more than the statement of Theorem 6.6. We have not only proved that the number of edges in such a bipartite graph is at most $a(n - a)$, with $a = \lfloor n/2 \rfloor$. We have also proved that this maximum can indeed be attained. In other words, the upper bound $a(n - a)$ cannot be further improved, since there are bipartite graphs on n vertices with exactly that many edges. In what follows, we will refer to this fact by saying that the upper bound $a(n - a)$ is *sharp*.

Third, we have also proved the fact that $a(n - a) \leq n^2/4$ is a direct consequence of the inequality between the *geometric* and *arithmetic mean*, that is, the inequality

$$\sqrt{xy} \leq \frac{x + y}{2}, \quad (6.1)$$

which holds for all nonnegative real numbers x and y . We will use this inequality many times. To prove this inequality, note that, after taking squares and rearranging, it becomes the inequality $(x - y)^2 \geq 0$. Because these steps are reversible, (6.1) follows. Now set $x = a$ and $y = n - a$ to get the inequality $\sqrt{a(n - a)} \leq n/2$.

Theorem 6.6 shows that, if a graph G on n vertices has more than $n^2/4$ edges, then G necessarily contains an *odd cycle* as a subgraph.

We claim that much more is true, namely, a graph with that many edges on that number of vertices will even contain a *triangle*—and therefore a very short odd cycle—as a subgraph.

Theorem 6.7 *If a graph G on n vertices contains more than $n^2/4$ edges, then it contains a triangle as a subgraph.*

We challenge the reader to prove this fact by induction, then check our solution in Exercise 1. An even stronger and more surprising fact (Supplementary Exercise 4) is that, if the graph G on $2n$ vertices has more than n^2 edges, so then G necessarily contains n triangles! Therefore, as long as G has

no more than n^2 edges, G may contain zero triangles, but as soon as G has more than n^2 edges, it suddenly contains at least n triangles.

Proof: (of Theorem 6.7) Let G be a graph on n vertices that contains no triangles. Let us try to tweak the proof of Theorem 6.6 so that it gives us this stronger conclusion. In that proof, we knew that our graph was bipartite, so it had two color classes with no edges within them, so the edges had to be between vertices of different color classes.

Here, we do not know that G is bipartite. However, we can still look at the largest *induced subgraph* of G that is empty and call its vertex set X . That is, there are no edges between two vertices of X , and there is no induced subgraph of G with more than $|X|$ vertices having that property.

Let Y be the set of vertices of G that are not in X . Then $|X| + |Y| = n$.

Crucially, no vertex of G can have degree more than $|X|$. Indeed, no two neighbors of any vertex can be adjacent, since that would mean that G contains a triangle. So the neighbors of any vertex form an empty induced subgraph of G , and as such, there cannot be more than $|X|$ of them.

While we do not know that G is bipartite, we do know that its edges are either within Y or between X and Y . This means that all edges of G have at least one vertex in Y . So if we add the degrees of the vertices in Y , we count each edge of G at least once. Therefore, the number $E(G)$ of edges of G satisfies

$$\begin{aligned} E(G) &\leq \sum_{y \in Y} d_y \leq \sum_{y \in Y} |X| = |X||Y| \\ &\leq \left(\frac{|X| + |Y|}{2} \right)^2 = \frac{n^2}{4}. \end{aligned}$$

Here the inequality between the two lines follows from (6.1), the inequality between the geometric and arithmetic mean. \diamond

Note that we used inequality (6.1) in the proof of Theorem 6.6 as well.

We point out that, in the above proof, in the step $E(G) \leq \sum_{y \in Y} d_y$, equality holds if and only if all edges of G have *exactly one* vertex in Y , that is, when G is bipartite with color classes X and Y . So it is again the complete bipartite graph that provides the highest number of edges without containing a triangle.

Now we know how many edges a graph on a given vertex set can have if it is not to contain a triangle as a subgraph. This question can be generalized in several natural directions. That is, a triangle is both a complete graph and a cycle. One can ask how many edges a graph on n vertices can have if it is not to contain the k -vertex complete graph K_k . One can also ask how many edges a graph on n vertices can have if it is not to contain the k -vertex cycle C_k . The latter can be further generalized by asking how many edges G can have if it is not to contain a cycle of length k or less.