Math 38 — Graph Theory Nadia Lafrenière

Basic definitions and some problems 03/28/2022

Can you draw these pictures, without ever crossing your path?

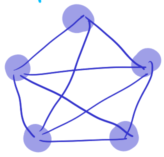 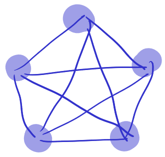

Can you draw this picture without ever lifting your pencil?

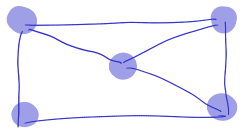

These are children problems, but also real-life problems in graph theory, namely to know whether a graph is planar, or similar to know if a graph is Eulerian.

The first problem: Seven bridges of Königsberg (Euler, 1736)

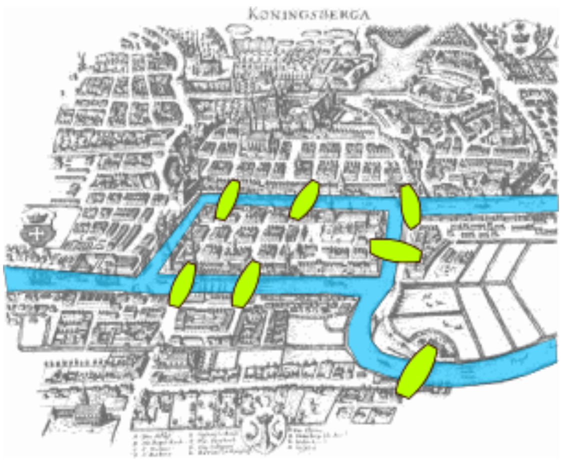

Euler was wondering if one can go from one place in the Königsberg area, and back to that original place, by taking every bridge exactly once.

(This is considered to be the first solved problem in graph theory).

A modelisation of the problem:

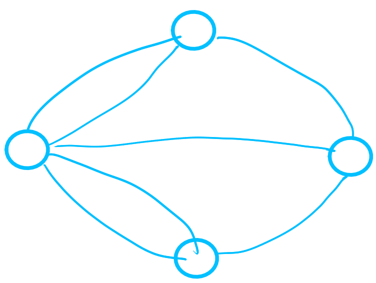

This graph model the areas of the city. There is no need to know the exact location of each bridge.

Remarks:

- Since we have to go back where we started, we do not care where we start.
- Everytime we go from a location to another and back, we cross 2 bridges adjacent to that location.

Since every island has an odd number of bridges, it is not possible to visit all the islands by taking every bridge exactly once.

Some definitions

A graph G is made of a set of vertices (modeling some objects), and a set of relations between two vertices, called the edges. We denote G = (V,E) for the graph with vertices V and edges E. Any edge is a pair of two vertices called the endpoints.

We draw a graph (on paper or on the computer) by representing the vertices as points, and we draw a curve between two vertices if they are endpoints of the same edge. We can draw differently the same graph.

Example

A loop is an edge whose endpoints are the same vertex. 

Multiple edges are edges having the same pair of endpoints.

A simple graph is a graph having no loop nor multiple edges.

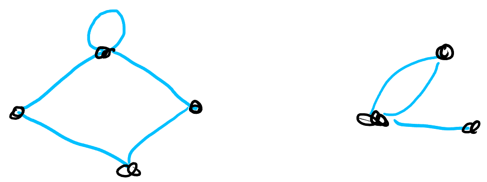

Not simple graphs

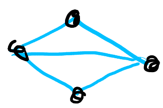

Simple graph

When uv (or equivalently) vu is an edge, we say the vertices u and v are adjacent, or that they are neighbors.

Subgraphs and containment

A graph $G'=(V',E')$ is a subgraph of $G=(V,E)$ if $V' \subseteq V$ and $E' \subseteq E$. We then say that G' is contained in G, denoted $G' \subseteq G$.

Example

Every graph with n vertices is a subgraph of the complete graph with $m \geq n$ vertices.

A graph is connected if, for every pair of vertices, there is a path (i.e. a sequence of edges) between them that belongs to the graph. It is otherwise disconnected.

Some important problems in graph theory

1. Acquaintances

Do every set of six people contain at least three mutual acquaintances or three mutual strangers?

That question can be represented using a graph. Every person is a vertex, and there is an edge between two persons if they know each other. Here, we assume knowing each other is a mutual relation, i.e. knowing a celebrity usually does not count.

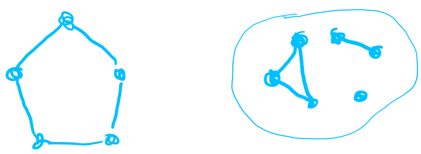

Two graphs. The first one is a 5-vertex graph with no three mutual strangers, nor three acquaintances. The second one has six vertices, and contain both three mutual strangers and three acquaintances (a clique).

As a homework, you will have to prove your solution to this statement.

Some useful vocabulary:

A clique in a graph is a set of pairwise adjacent vertices, i.e. a complete subgraph.

An independent set is a subset of vertices with no adjacent pairs.

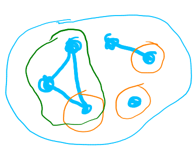

- A clique
- An independent set

2. Job assignments

If there are m jobs and n people, not all qualified for all the jobs, is there a way we can fill all the jobs?

Definition

A bipartite graph is the disjoint union of two independent sets.

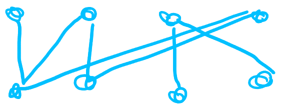

The edges are between a job and a qualified person for that job.

(The jobs cannot all be filled in this example).

3. Scheduling and avoiding conflicts

My high school used to have a very long exam sessions at the end of the year, and there were still some conflicts. I wish the administrators knew graph theory...

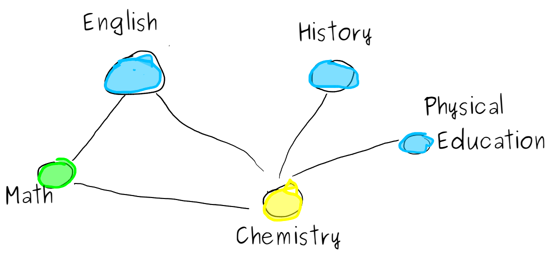

Vertices: Subjects

Edges: If someone takes both subjects, i.e. eventual scheduling conflicts.

A coloring of a graph is a partition of a set into independent sets. Scheduling with no conflicts is equivalent to coloring. If we want to use the minimum time, we should use as few colors as possible.

Schedule:

1. History—English—PE
2. Chemistry
3. Math

Reference: Douglas B. West. Introduction to graph theory, 2nd edition, 2001. Sections 1.1.1 and 1.1.2.

Math 38 — Graph Theory

Nadia Lafrenière

Matrices for graph and Isomorphisms

03/30/2022

We saw last class that two graphs are the same if they are differently, as long as we are simply "moving the vertices". The goal of today's lecture is to make this statement more formal. One tool we will use is adjacency and incidence matrices. We will as well start classifying the graphs.

Matrices: adjacency matrix and incidence matrix

Let $G=(V, E)$ be a graph without any loop (it does not have to be a simple graph). We number the vertices from 1 to n and the edges from 1 to m.

The adjacency matrix of G, written $A(G)$, is the matrix whose (i,j)-entry is the number of edges with endpoints the vertices i and j.
The incidence matrix of G, written $M(G)$, is the n-by-m matrix whose (i,j)-entry is 1 if vertex i is an endpoint of edge j, and otherwise 0.
The adjacency matrix is always a symmetric matrix.
The graph on the left has the following adjacency and incidence matrices:

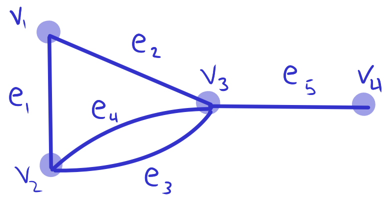

$$A(G) = \begin{pmatrix} 0 & 1 & 1 & 0 \\ 1 & 0 & 2 & 0 \\ 1 & 2 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \qquad M(G) = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 1 & 0 \\ 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}$$

The degree of a vertex (in a loopless graph) is the number of edges incident to that vertex.

Isomorphisms

So when are two graphs the same? We will answer this question using the notion of a bijection. As a reminder, this an injective and surjective function, or a one-to-one correspondence.

An isomorphism from a simple graph G to a simple graph H is a bijection $f:V(G)\to V(H)$ such that every edge uv of G is mapped to the edge $f(u)f(v)$ of H. We then say G and H are isomorphic, denoted $G \cong H$.

This is equivalent to asking that there exists a simultaneous permutation of the rows and columns of the adjacency matrix of G that would yield the adjacency matrix of H.

Example

The following graphs are isomorphic:

This is easily seen with the bijection that exchanges 1 and 3.

Remarks:

- Finding a bijection of the labels is the way to prove two graphs are isomorphic. However, to prove they are not isomorphic, there are many ways. For example, if the list of degrees is not the same, you will never be able to find an isomorphism. Or if the number of edges (or edges) do not correspond. Among others.
- The isomorphism relation is an equivalence relation, i.e. this is a symmetric relation ($G \cong H$ iff $H \cong G$), a transitive relation ($G \cong H$ and $H \cong J$ imply $G \cong J$) and a reflexive one ($G \cong G$). That means that we can split the graph into equivalence classes.

Example

The following graphs are not isomorphic. They both have six vertices, all of degree 3, and nine edges, and they are both connected, but one is bipartite and the other is not. Since they don't have the same properties, they are not isomorphic.

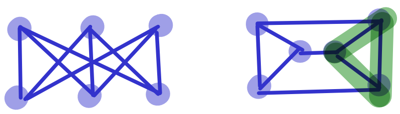

No triangle appear in the first graph.

Example

All the isomorphism classes for graphs with 4 vertices are

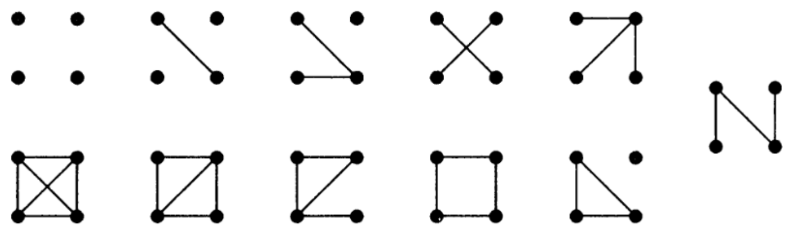

Special graphs

There are some graphs that have special names, and that turns out to be handy for whenever we want to use them or to classify them.

Complete graphs: Graphs with n vertices and $\binom{n}{2}$ edges.

K_n Example: K_5 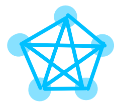

Complete bipartite graphs: Bipartite graphs with independent sets of size s and r, with sr edges.

$K_{s,r}$ Example: $K_{4,2}$ 

Paths: Connected graphs, with all the vertices of degree 2, except at most two who have degree 1.

P_n Example: P_4 

Cycles: Paths with as many edges as vertices.

C_n Example: C_5 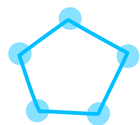

The complement of the graph G is the graph that has the same vertices and whose edges are all the edges that do not belong to G: $K_{|V|} - E(G) = \overline{G}$.

A graph G is self-complementary if its complement $\overline{G}$ is isomorphic to G.

Example: C_5 is self-complementary.

$K_5 =$ 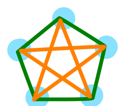  $\cong$ 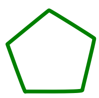

A decomposition of a graph is a list of subgraphs in which every edge appears exactly once.

Example: The cube decomposed into copies of $K_{1,3}$

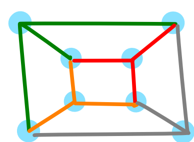

Note: $K_{1,3}$ is often called the claw.

Proposition

A graph G is self-complementary if and only if the complete graph is a decomposition into two copies of G.

The Petersen graph

The Petersen graph is a 10-vertices graph with 15 edges that is very famous, as it is an example or a counter-example to many phenomena.

The Petersen graph is the graph of 2-element subsets of {1,2,3,4,5}, and there is an edge between 2 subsets if their intersection is empty.

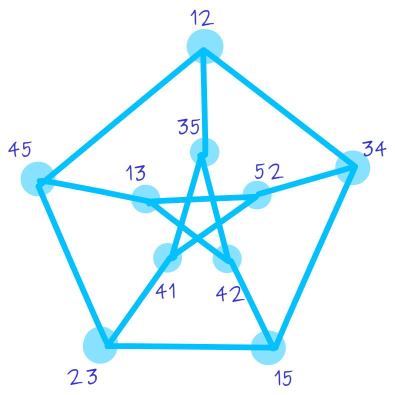

Some properties of the Petersen graph:
- Two non-adjacent vertices share exactly one neighbor.
- The graph has no triangle, but is not bipartite.
- The shortest cycle in the Petersen graph has length 5. (The length of the shortest cycle in a graph is called the girth of the graph.)

Reference: Douglas B. West. Introduction to graph theory, 2nd edition, 2001. Section 1.1

Math 38 – Graph Theory — Nadia Lafrenière

Connection in graphs — 04/01/2022

Today's lecture aims to define the proper vocabulary to talk about trajectories and connectedness in graphs.

Definitions

Recall that a path is a graph whose vertices can be ordered without repetition (except maybe for the endpoints) in a sequence such that two consecutive vertices are adjacent. A path is a u,v–path if it starts at vertex u and ends at vertex v.

A walk is a list $(v_0, e_1, v_1, \ldots, e_k, v_k)$ of vertices and edges such that the edge e_i has endpoints v_{i-1} and v_i. A walk is a u,v–walk if its endpoints (the first and last vertices of the walk) are u and v. If there is no multiple edges, we can write the walk as $(v_0, v_1, \ldots, v_k)$.

A trail is a walk with no repeated edge. Similarly, a u,v–trail has endpoints u and v.

The points that are not endpoints are internal vertices.

The length of a walk, trail, path or cycle is its number of edges.
A walk or a trail is closed if its endpoints are the same.

Example

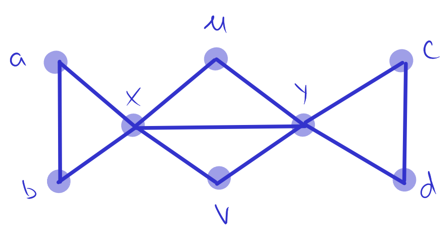

(a,x,a,b,x,u,y,x,a) specifies a closed walk, but not a trail (ax is used more than once).

(a,b,x,u,y,x,a) specifies a closed trail.

The graph contains the five cycles (a,b,x,a), (u,y,x,u), (v,y,x,v), (x,u,y,v,x) and (y,c,d,y).

The trail (x,u,y,c,d,y,v,x) is not an example of a cycle, since vertex y is repeated (so it is not a path).

Lemma

Every u,v-walk contains a u,v-path.

Proof

The proof can be done using the principle of strong induction, and we induce on the number of edges.

Base case: No edge, u=v is the only vertex in the graph. Only walk has length 0, and is therefore a path.

Induction hypothesis: Assume that, for a walk with $k<n$ edges, there is always a path with the same endpoints.

Induction step: The walk has n edges. There are two cases: either there is no repeated vertex or only the endpoint is repeated, and then the walk is already a path, or there is a repeated vertex x. In the latter case, we delete the edges between the first and last occurrences of x, which leaves us with only one copy of x, and a u,v-walk with fewer than n edges. We can thus use the induction hypothesis to conclude that there exists a u,v-path in the u,v-walk.

■

Example: The u,v-walk from previous page.

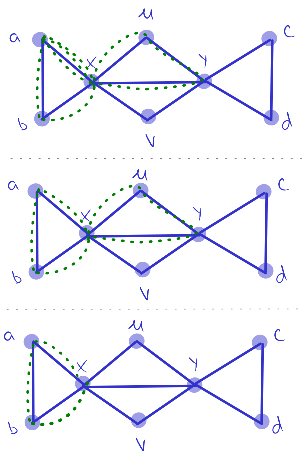

In the walk (a,x,a,b,x,u,y,x,a), we delete what happens between the first two occurrences of a, and get the closed walk (a,b,x,u,y,x,a). Then we delete what happens between the two occurrences of x, and get the cycle (a,b,x,a), which is a path.

Connectedness, components and cuts

Recall that a graph is connected if and only if there exists a path between u and v for every pair of vertices {u,v}.

A component of a graph G is a maximal connected subgraph.
A component is trivial if it has no edges; in this case, the unique vertex is said to be an isolated vertex.

Example

The following graph has 4 components, each of which are circled in orange.

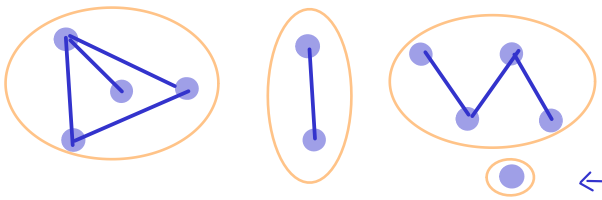

Proposition

Every graph with n vertices and k edges has at least $n-k$ components.

Proof

The proof can be done by induction on k. The case of $k>n$ is obvious, since the number of components is always nonnegative.

Base case: If $k=0$, then each of the n vertices are isolated, and there are n components.

Induction hypothesis: Assume that a graph with $k-1$ edges and n vertices has at least $n-k+1$ components.

Induction step: Let $G=(V,E)$ with $|V|=n$ and $|E|=k$. Remove the edge e to get $G-e$. The component of G containing e can either be split into two components by removing e, or stay a component. So G has either the same number of components as $G-e$, or one fewer. By induction hypothesis, $G-e$ has at least $n-k+1$ components, so G has at least $n-k$.

∎

In the last proof, we had to distinguish the cases where removing the edge was creating a new component or not. An edge whose deletion creates new component has a special name:

A cut-edge or cut-vertex of a graph is an edge or vertex whose deletion increases the number of components. We write $G-e$ or $G-M$ for the subgraph of G obtained by deleting an edge e or a set of edges M; we write $G-v$ and $G-S$ for the graph obtained by deleting a vertex v or a set of vertices S along with their incident edges.

A subgraph obtained by deleting a subset of vertices and their incident edges is an induced subgraph: we denote it G[T] if T=V\S and we deleted the vertices in S.

Example

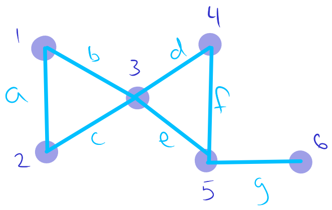

Vertices 3 and 5 are cut-vertices, and the edge g is the only cut-edge.
The induced subgraph for the vertices 1, 2, 3, 4 and 6:

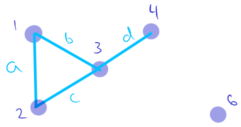

Theorem

An edge is a cut-edge if and only it if belongs to no cycle.

Proof

Let e=uv be an edge in the graph G, and let H be the component containing e. We can restrict the proof to H, since deleting e does not influence the other components. We want to prove that H−e is connected if and only if e is in a cycle in H.

If e is in a cycle c, c−e is a path P between v and u avoiding the edge e. To show that H−e is still connected, we need to show that, for every pair of vertices {x,y}, there is a path between x and y. Since H is connected, there exists in H such a path. If that path does not contain e, it is still in H−e. Otherwise, replace e by P, and remove an edge from that path everytime it appears twice consecutively.

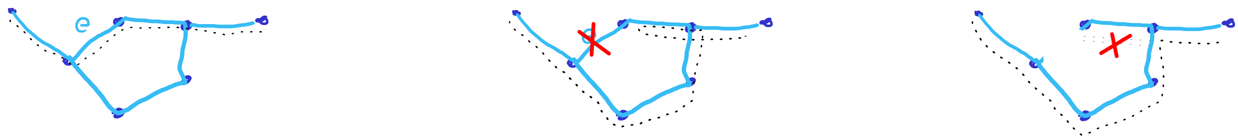

If H−e is connected, then there exists in it a path P between u and v. Hence, adding edge e=uv creates the cycle P+e. ∎

The last theorem allows us to characterize cut-edges. Would such a theorem be possible for cut-vertices? The following example proves that asking for it to be outside a cycle is not a requirement for a cut-vertex, since vertex 3 is a cut-vertex, and belongs to two cycles:

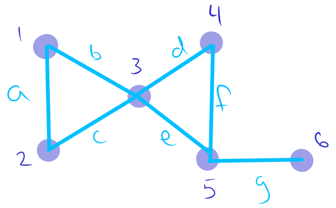

Removing vertex 3:

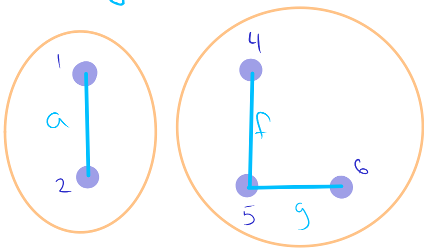

Two connected components

Reference: Douglas B. West. Introduction to graph theory, 2nd edition, 2001. Section 1.2

Math 38 — Graph Theory Nadia Lafrenière

Bipartite and Eulerian Graphs 04/04/2022

Today's lecture aims to give the important properties of bipartite graphs. We will also define Eulerian circuits and Eulerian graphs: this will be a generalization of the Königsberg bridges problem.

Characterization of bipartite graphs

The goal of this part is to give an easy test to determine if a graph is bipartite using the notion of cycles: König theorem says that a graph is bipartite if and only if it has no odd cycle.

Lemma

Every closed walk of odd length contains an odd cycle. This is called an odd closed walk.

Proof

We prove it using strong induction on the length of the walk (i.e. the number of edges).

Base case: length 1. The walk is a loop, which is an odd cycle.

Induction hypothesis: If a walk has odd length at most n, then it contains an odd cycle.

Induction step: Consider a closed walk of odd length n+1. If it has no repeated vertex (except the first and last one), this is a cycle of odd length. Otherwise, assume vertex v is repeated. We can split the walk into two closed walks starting and ending at v, one of even length, and one of odd length smaller than n. By induction hypothesis, the latter contains an odd cycle.

□

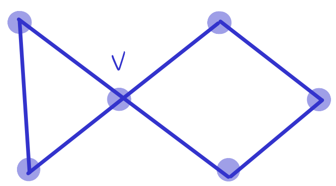

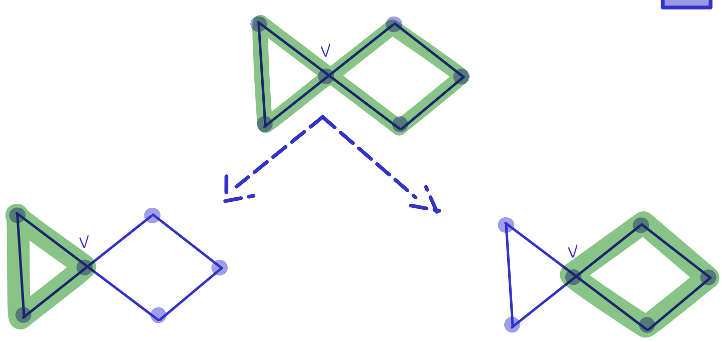

That lemma will be helpful for characterizing bipartite graphs. Of course, bipartite graphs can have even cycles, which starts in one independent set and ends there.

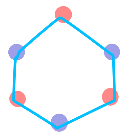

We can represent the independent sets using colors.

Theorem (König, 1936)
A graph is bipartite if and only if it has no odd cycle.

Proof
Notice that a graph is bipartite if and only if all its components are bipartite. So we do the proof on the components.
⇒ We prove the contrapositive: it is has an odd cycle, it is not bipartite.
Since every cycle must end at the vertex where it starts, it starts and ends in the same independent set. Since every edge is going from one set to the other, we alternate between the two sets. At the end of the cycle, we cannot close it, since we would need to change the set of the first vertex. Hence, if a connected graph is bipartite, it has no odd cycle.

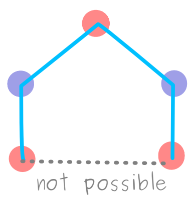

⇐ We still need to prove that a connected graph without odd cycle is bipartite. If the graph has only one vertex, it is bipartite. Otherwise, start at vertex u, and color its neighbors with color blue. Then, color the neighbors of the blue vertices in red, and repeat this process by coloring the neighbors of the red vertices in blue, until all vertices have been colored. I claim that no vertex will change color in that process; assume otherwise, that v is changing color. That would mean that there exists a path of odd length from u to v (the one that colors v in blue), and a path of even length doing it (the one that colors v in red). The combination of these two paths is an odd walk, and contains an odd cycle, which is prohibited by the hypothesis. Hence, the coloring is well defined, and the two colors represent independent sets. The graph is bipartite. ∎

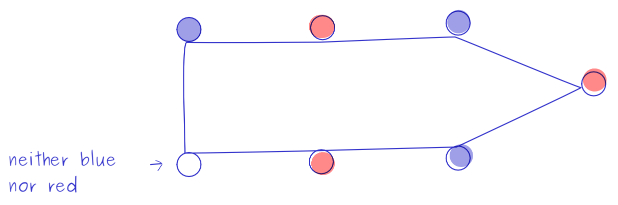

odd cycle

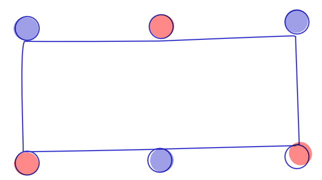
even cycle

Technique for checking whenever a graph is bipartite:
- If it is bipartite, prove it by finding two independent sets.
- If it is not bipartite, find an odd cycle.

Eulerian circuits

A graph is Eulerian if it has a closed trail containing all the edges.

The graph in the Königsberg bridges problem is not Eulerian. We saw that the fact that some vertices had odd degree was a problem, since we could never return to that vertex after leaving it for the last time.

Theorem
A graph is Eulerian if and only if it has at most one nontrivial component (i.e. component with edges), and if every vertex has even degree.

Proof
We first prove ⇒ by proving the contrapositive: if a graph has more than one non-trivial component, or it there is a vertex of odd degree, then the graph is not Eulerian.

If a graph has at least two non-trivial components, there can't be a walk going through all the edges, since they are in separate components.

If a graph has a vertex of odd degree, we are in the case of the Königsberg bridges: we can leave the vertex more often than we can come back (or vice-versa), and thus our trail cannot be closed.

⇐ We need to prove that a connected graph with only vertices of even degrees is Eulerian. We can ignore the isolated vertices for this since we are focusing on the edges. The following lemma is useful:

Lemma

If every vertex of a graph has degree at least 2, then it contains a cycle.

Proof

Let P be a maximal path in that graph. If it is a cycle, we are done. Otherwise, let u be an endpoint of P. Since it has degree at least 2, u has a neighbor v not in P. But since P is maximal, that means that v is already in P, and the edge uv completes the cycle.

Proof of the theorem (continued)

We proceed by induction on the number of edges.
Base case: 0 edge, the graph is Eulerian.
Induction hypothesis: A graph with at most n edges is Eulerian.
Induction step: If all vertices have degree 2, the graph is a cycle (by definition) and it is Eulerian. Otherwise, let G' be the graph obtained by deleting a cycle. The lemma we just proved shows it is always possible to delete a cycle. By induction hypothesis, G' is Eulerian. To build an Eulerian circuit in G, start by the cycle we just deleted, and append the Eulerian circuit of G'. ∎

Proposition

Every graph with only vertices of even degree decomposes into cycles.

Eulerian circuits are closed trails that pass through all edges. A similar property is being Hamiltonian: a Hamiltonian circuit is a circuit that passes though all vertices exactly once. A Hamiltonian graph is a graph with a Hamiltonian circuit.

Reference: Douglas B. West. Introduction to graph theory, 2nd edition, 2001. Section 1.2

Math 38 – Graph Theory Nadia Lafrenière

Vertex Degrees and counting 04/06/2022

Today, we are doing a bit of combinatorics and will deduce some properties on the degrees, number of edges and number of vertices.

We already defined the degree of a vertex in a loopless graph to be the number of edges incident to it.

For a general graph, define the degree $d_G(v)$ of the vertex v to be the number of edges incident to it, with each loop counted twice.

The maximum degree of a vertex is denoted $\Delta(G)$ and the minumum degree is denoted $\delta(G)$.

A graph is said to be regular if $\delta(G) = \Delta(G)$.

The order of a graph $G=(V,E)$ is $|V|$, as the size of G is $|E|$.

Example

- K_n is a regular graph. Each vertex has degree $n-1$.
- $K_{m,n}$ is regular if and only if $m=n$. Then, the degree is always n.
- A connected regular graph that has the same order and size is a cycle.
- Hypercubes are regular graphs.

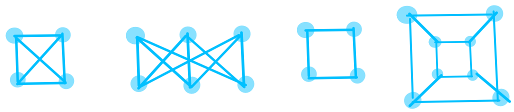

Counting and bijections

Proposition (degree-sum formula)

If $G=(V,E)$ is a graph, then $\sum_{v \in V} d_G(v) = 2|E|$.

Proof

For each edge, there are two endpoints (maybe equal). If the endpoints are different, this edge contributes for 1 in the degree of two different vertices. If the edge is a loop, it adds 2 to the degree of the vertex it is incident to. So either way, every edge accounts for 2 in the total degree count. ■

Corollary

In any graph $G=(V,E)$, the average degree is $2|E|/|V|$, and $\delta(G) \le 2|E|/|V| \le \Delta(G)$.

Corollary

Every graph has an even number of vertices of odd degree.

Corollary

A k-regular graph (i.e. a regular graph in which the degree of each vertex is k) has $k|V|/2$ edges.

Example: Hypercubes

The n-dimensional hypercube H_n is defined recursively as:

- H_0 is the simple graph with one vertex
- H_{n+1} is obtained by creating two copies of H_n and appending an edge between corresponding vertices in the two copies.

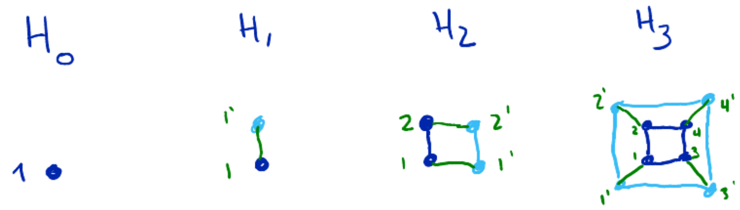

Proposition

H_n is regular, as each vertex has degree n.

Proof

The proof can be done by regular induction.
The base case is H_0, and it has no edge.
Induction hypothesis: The n-dimensional hypercube H_n is n-regular.
Induction step: The (n+1)-dimensional hypercube is made of two copies of H_n, and we add an edge between every pair of similar vertices in the two copies. This way, we add exactly one to the degree of each vertex from H_n, and that degree is, by induction hypothesis n. ■

Proposition

If k>0, then a k-regular bipartite graph has the same number of vertices in its two independent sets.

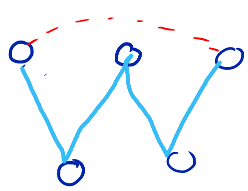

Either not bipartite or not regular.

Proof

Since the graph is regular, all vertices have degree k. If there are m edges in total, the sum of the degrees for all the vertices in one independent set is m, as every edge has exactly one endpoint in that set. Since the graph is k-regular, there are m/k vertices in each set, so the order of both sets is the same.

■

Vertex-deletion and reconstruction conjecture

Is it possible to reconstruct a graph if you have only a list of its subgraphs? There is a long-standing, and still open conjecture saying it is, and so far we know it is almost always possible (that being understood in a probabilistic sense).

For a graph G, a vertex-deleted subgraph is an induced subgraph $G-v$ obtained by deleting a single vertex v.

Example

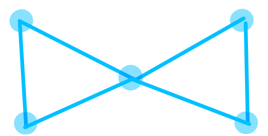

has vertex-deleted subgraphs 4 x

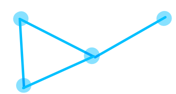

1 x

Proposition

For a simple graph $G=(V,E)$ of order $n>2$ and size m,

$$m = \frac{\sum_{v\in V} \#E(G-v)}{n-2}$$

where $\#E(G-v)$ is the number of edges in the graph $G-v$.

Proof

We start with the summation, and we will prove the summation is equal to $m(n-2)$:

$$\sum_{v\in V} \#E(G-v) = \sum_{v\in V} |E| - d_G(v) = \sum_{v\in V} |E| - \sum_{v\in V} d_G(v) = mn - 2m$$

Conjecture (Reconstruction Conjecture – Kelly, Ulam, 1942)

If G is a simple graph with at least three vertices, then G is uniquely determined by the list of its vertex-deleted subgraphs (up to isomorphism).

Note that the hypothesis that G has at least three vertices is important. Otherwise, we would find a counterexample with two vertices, since both simple graphs with two vertices have the same set of vertex-deleted subgraphs.

Example

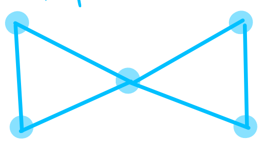

has vertex-deleted subgraphs 4 x

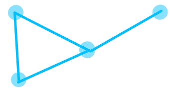

1 x

To reconstruct the graph, we know that 4 vertices have degree #E(G)−4 and 1 has degree #E(G)−2. Using the proposition, the number of edges in G is (2+4x4)/3=6. So the list of degrees is (2,2,2,2,4), and the graph is connected.
That means that the vertices are in two cycles. The length of the cycles can be found by looking at the subgraphs: there is at least one cycle of length 3. Since the graph is simple, both cycles have length 3 and the graph has to be isomorphic to the bowtie.

Even though the conjecture is not proven, there are a number of cases that are known. Also, we can know some properties from the list of subgraphs; for example if the graph is connected.

Reference: Douglas B. West. Introduction to graph theory, 2nd edition, 2001. Section 1.3

Math 38 — Graph Theory Nadia Lafrenière

Extremal problems 04/08/2022

Extremal problems consider the minimum and maximum numbers some statistics on a class of graphs can reach. We introduce some of the types of proofs useful in graph theory: Algorithmic, and by construction.

First example

In any simple graph (V,E), the maximum number of edges is

$$\binom{|V|}{2} = \frac{|V|(|V|-1)}{2}$$

Proof

In a simple graph, there can be at most one edge per pair of distinct vertices. The maximum number of edges appear in $K_{|V|}$.

This is an extremal problem, since we are looking at the maximum number of edges. The class of graphs here is all simple graphs.

Example

In a bipartite graph with independent sets of size k and m, there can be at most km edges.

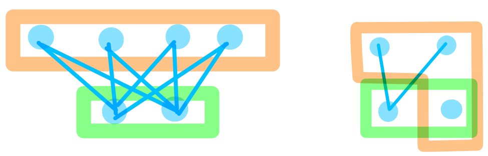

Independent sets of size 2 and 4, 8 edges at maximum. km is the number of edges of $K_{m,k}$.

Edges in connected graph

Proposition

The minimum number of edges in a connected graph with n vertices is n−1.

Proof

We need to prove two things:

- If a graph with n vertices has fewer than n−1 edges, it is not connected.
- There exists a connected graph with n vertices and n−1 edges.

Recall from last week (Friday), that a graph with n vertices and m edges has at least n−m components. Hence, if m<n−1, the graph has at least 2 components and is not connected.
Also, the path with n vertices has n−1 edges and is connected, proving that the minimum is realized.

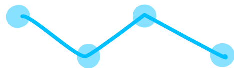

∎

Remark (on the proof technique)
When giving the solution to an extremal problem, there are two parts to be proven:
- That the value we give is minimal (or maximal), i.e. that you cannot give a lower (respectively, higher) value.
- That this value can be realized on at least one graph of the class we consider.

Proposition
Let G be a simple graph with n vertices. If the minimum degree is $\delta(G) \geq (n-1)/2$, G is connected.

Proof
The minimum degree of the graph means that every vertex should have at least this number of neighbors, in a simple graph.
To prove that G is connected, we must show that there is a path between any pair of vertices $\{u,v\}$. We will in fact prove that there exists a path of length at most 2.
- If $\{u,v\}$ are adjacent, they are obviously in the same component.
- Otherwise, they share at least one neighbor w: There are n−2 other vertices, and the sum of their degree is $d(u)+d(v) \geq n-1$. Hence, u−w−v is a path connecting them.

∎

A bound is said to be sharp if improving it (reducing a lower bound or increasing an upper bound) would make the statement wrong.
The bound in the last problem is sharp. To prove it, we give an example of a graph with n vertices and minimum degree $\lfloor \frac{n}{2} \rfloor - 1$ that is not connected: This graph is the disjoint union of $K_{\lfloor \frac{n}{2} \rfloor}$ and $K_{\lceil \frac{n}{2} \rceil}$.

K_5, degree 4

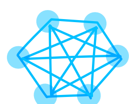

K_6, degree 5

11 vertices
Minimum degree is 4, just under 5 = (11−1)/2.
Graph is disconnected.

Bipartite subgraph

Here we prove that, given a graph G, we can always find a bipartite subgraph with at least a fixed number of edges. We give an algorithmic proof to construct the graph, but a proof can also be done by induction.

Theorem

Every loopless graph $G=(V,E)$ has a bipartite subgraph with at least $|E|/2$ edges.

Proof (algorithmic)

We start with any partition of the vertices into two sets X and Y. Let H be the subgraph containing all the vertices, but only the edges with one endpoint in X and one in Y.

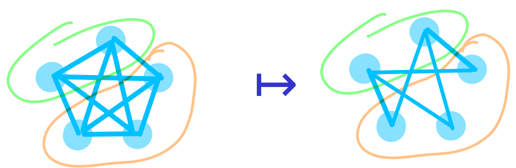

6 edges, instead of 10

Let v be a vertex in X. If H has fewer than half the edges incident to v, then it means that v has (in G) more neighbors in X than in Y. To increase the number of edges in H, switch v to Y. The number of edges just increased.

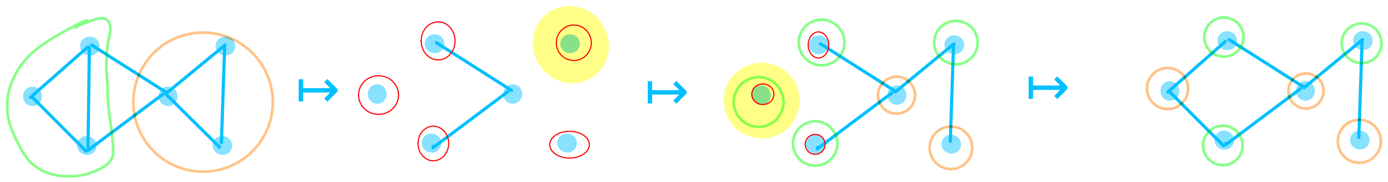

○ = less than half the edges

As long as H does not have at least half the edges of G at every vertex, there are vertices that can be swapped from X to Y or Y to X; repeat this process. When it terminates, the number of edges in H is always at least half the number of edges of G. ∎

Triangle-free graphs

A graph is said to be triangle-free if it has no three vertices that are all adjacent. In general, a graph G is H-free if it does not contain H as a subgraph.

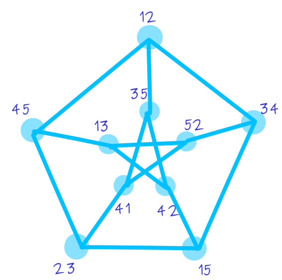

The Petersen graph is triangle-free (but not bipartite).

Theorem (Mantel, 1907)

The maximum number of edges in a simple triangle-free graph with n vertices is $\lfloor \frac{n^2}{4} \rfloor$.

Proof

For the proof, we again need to prove two things:
- that a triangle-free graph with n vertices cannot have more than $\lfloor \frac{n^2}{4} \rfloor$ edges.
- that there exists, for any n, a graph with n vertices and $\lfloor \frac{n^2}{4} \rfloor$ edges that has no triangle.

For the first part, assume the graph is triangle-free. Take a vertex v of maximal degree Δ. Its Δ neighbors cannot have edges among them. So every edge of G must have at least one endpoint in a non-neighbor of v, or in v itself. There are n−Δ such vertices. Each such vertex has degree as most Δ.

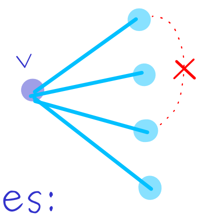

Therefore, we give an upper bound on the number of edges: the number of edges is at most Δ(n−Δ) (because n−Δ is the number of vertices not adjacent to v). Maximizing Δ(n−Δ) gives Δ=n/2. Hence, the number of edges is at most $\lfloor \frac{n^2}{4} \rfloor$.

For the second part, we must prove that a triangle-free graph has $\lfloor \frac{n^2}{4} \rfloor$ edges. This is the case of $K_{\lceil \frac{n}{2} \rceil, \lfloor \frac{n}{2} \rfloor}$.

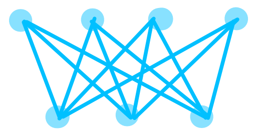

We can split 7 vertices into two sets of 3 and 4 vertices, which leads to 12 edges:, which is the smallest integer below 49/4.

Reference: Douglas B. West. Introduction to graph theory, 2nd edition, 2001. Section 1.3