Graph Theory

Jan-Christoph Schlage-Puchta

Graph Theory

Jan-Christoph Schlage-Puchta

1 Graphs

Informally a graph consists of a set of points, called vertices, some of which are connected by edges. Formally a graph $\mathcal{G} = (V, E)$ consists of a set V of vertices and a set E of edges, where $E \subseteq V^{(2)} = \{\{v, w\} | v, w \in V, v \neq w\}$. Graphs can be used to model relations both among mathematical and real world objects. Depending on the problem to be modeled the notion of a graph can be varied: A graph may have loops, i.e. edges connecting a vertex to itself, multiple edges, i.e. between two points there may be more than one edge, or directed, i.e. each edge is not a 2-element set but an ordered pair. A path in a graph is a sequence of vertices $v_1, \dots, v_n$, such that each vertex is connected to the following one by an edge.

One way of encoding graphs is the so called adjacency matrix. Assume that V is finite. Without loss we may put $V = \{1, 2, \dots, n\}$. Then define a matrix A by $A = (a_{ij})_{i,j=1}^n$, where $a_{ij} = \begin{cases} 1, & \{i, j\} \in E \\ 0, & \{ij\} \notin E \end{cases}$. Then A is a symmetric 0-1-matrix with all diagonal elements equal to 0. If we consider graphs with loops, diagonal elements may be non-zero, in multiple graphs entries are arbitrary non-zero integers, and in directed graphs this matrix would not be symmetric.

The adjacency matrix encodes interesting information on the graph.

Proposition 1. *Let $A = (a_{ij})_{i,j=1}^n$ be the adjacency matrix of a graph $\mathcal{G}$. Then $A^k = (b_{ij}^{(k)})_{i,j=1}^n$, where $b_{ij}^{(k)}$ equals the number of different paths from i to j in $\mathcal{G}$ of length k.*

Proof. Proof by induction over k. The case $k = 1$ is trivial.

We cut a path p from i to j of length k into an initial segment p' of length $k - 1$, and the final step. Let ν be the last point of p'. Then $\{\nu, j\} \in E$, since we can reach j from ν. Moreover, paths having different vertices at their next to last position are different, and paths with the same next to last vertex are different if and only if their initial segments are different. Hence

$$b_{ij}^{k)} = \sum_{\substack{\nu \in V \\ \{\nu, j\} \in E}} b_{i\nu}^{(k-1)} = \sum_{\substack{\nu \in V \\ a_{\nu j}=1}} b_{i\nu}^{(k-1)} = \sum_{\nu=1}^{n} a_{\nu j} b_{i\nu}^{(k-1)}.$$

But the last expression is exactly the (i,j)-entry of $A \cdot (b_{ij}^{(k)})_{i,j=1}^{n}$, which by the inductive hypothesis equals $AA^{k-1} = A^k$. $\square$

For example, the trace of A^3 equals the number of triangles in $\mathcal{G}$.

Exercise. Formulate and prove a version of Proposition 1 for directed graphs and graphs with multiple edges!

A graph is called connected, if any two vertices are connected by a path. It follows from Proposition 1 that $\mathcal{G}$ is connected if and only if there exists some n, such that all entries of A^n are positive. Computing A^n for large n is best done by diagonalizing A: Find a mtarix B, such that BAB^{-1} is diagonal. Then $A^n = B^{-1}(BAB^{-1})^n B$, and the inner power can be computed by taking powers of complex numbers. Hence knowing the eigenvalues of A should be useful.

Two graphs (V, E), (V', E') are called isomorphic, if there exists a bijection $\varphi : V \to V'$, such that $\{v, w\} \in E$ if and only if $\{\varphi(v), \varphi(w)\} \in E'$. For all practical purposes isomorphic graphs can be considered as equal.

The question whether two graphs are isomorphic or not can become quite difficult. Therefore it is useful to have practical criteria which sole this problem in most cases.

Proposition 2. *Two graphs with adjacency matrix A, A', respectively, are isomorphic if and only if there exist a permutation matrix P, such that $A' = PAP^{-1}$.*

Proof. Conjugating by a permutation matrix is equivalent to permuting rows and columns. This permutation in turn corresponds to an isomorphism. $\square$

The easier question, whether for two matrices A, A' there exists an invertible matrix B, such that $A' = BAB^{-1}$ is solved in linear algebra: Such a matrix exists if and only if A and A' have the same Jordan normal form. Since the adjacency matrix of a graph is symmetric, the problem becomes even easier: Two adjacency matrices are conjugate if and only if they have the same eigenvalues with the same multiplicities. Unfortunately it is not clear how to pass from the property ”conjugate” to ”conjugate by permutation matrices”, so two graphs having the same eigenvalues may still be non-isomorphic. Still, comparing the eigenvalues is a quick way to check whether graphs could be isomorphic or not.

Exercise. Let $\mathcal{G}, \mathcal{G}'$ be graphs with adjacency matrices A, A'. Prove that if A and A' are conjugate, then $\mathcal{G}$ and $\mathcal{G}'$ have the same number of vertices, and the same number of triangles. More generally show that for every integer k the number of paths $(v_1, v_2, ldots, v_k$ with $v_i = v_k$ are the same in $\mathcal{G}$ and $\mathcal{G}'$.

2 Spectral theory of graphs

Let $\mathcal{G}$ be a graph, A the adjacency matrix of $\mathcal{G}$. The set of eigenvalues of $\mathcal{G}$ is called the spectrum of $\mathcal{G}$.

Eigenvalues of $\mathcal{G}$ have a physical interpretation. Imagine that each vertex is a small mass, and each edge is a spring connecting two masses. If we pull some of the masses in one direction, and others in the opposite direction, and let go, the whole system starts to oscillate. Just as in the continuous case, eigenvalues correspond to simple modes of oscillation. For example, if we consider a circle of $4n$ points, we could have either a rather slow oscillation around one axis of symmetry, or a rather fast oscillation where all even numbered points remain fixed, all points with number $\equiv 1 \pmod 4$ move synchronously, and all points with number $\equiv 3 \pmod 4$ move synchronously, but in opposite direction as the points $\equiv 1 \pmod 4$. The physical reason for the difference in the frequency is that in the first case we have that all points move pretty much like their neighbours, thus the springs transmit little energy, while in the second case each point behaves quite different from its neighbours, thus the springs tramsit lots of energy. We conclude that the size of eigenvalues, which gives the speed of oscillation, is related to the structure of the corresponding eigenfunctions by measuring how much this function differs on one pointfrom its neighbours.

This physical intuition can be made more precise. For the sake of simplicity we assume that $\mathcal{G}$ is a k-regular graph, that is, each vertex is contained in exactly k edges. We say that a graph $\mathcal{G}$ is bipartite, if the vertex set V can be partitioned into two sets A, B, such that all edges in $\mathcal{G}$ have one vertex in A and one vertex in B.

Proposition 3. *Let $\mathcal{G}$ be a k-regular graph.*

1. *The vector $(1, 1, \dots, 1)$ is an eigenvector to the eigenvalue k.*
2. *$\mathcal{G}$ is connected if and only if the eigenvalue k is simple.*
3. *If $\mathcal{G}$ is connected, then $\mathcal{G}$ is bipartite if and only if $-k$ is an eigenvalue of k.*

Proof. We only show part $\Leftarrow$ of (3), all other statements are much simpler and left to the reader. Let $f : V \to \mathbb{C}$ be an eigenfunction to the eigenvalue $-k$. Let v be a vertex such that $|f(v)|$ is maximal. Then we have

$$-kf(v) = Af(v) = \sum_{w\sim v} f(w) \geq k \min_{w\sim v} f(w) \geq k \min_{w\in V} f(w) \geq -k \max_{w\in V} |f(w)| = -k|f(v)|,$$

that is, we have equality throughout. In particular for all vertices adjacent to v we have $f(w) = -f(v)$. But then $|f(w)|$ is maximal for all vertices w adjacent to v, that is, since $\mathcal{G}$ is connected, that f takes only the values k and $-k$. Define A to be the set of vertices satisfying $f(v) = k$, and B the set of vertices satisfying $f(v) = -k$. Then A and B are disjoint, their union equals V, and a vertex in A is only connected to vertices in B. Hence $\mathcal{G}$ is bipartite. $\square$

Exercise. Prove the remaining statements of Proposition 3.

Let $\mathcal{G} = (V, E)$ be a graphs, and let $\mathcal{F}$ be the set of functions $f : V \to \mathbb{C}$, and define a scalar product on L^2 by putting $\langle f, g\rangle = \sum_{v\in V} f(v)\overline{g(v)}$, and the L^2-norm on $\mathcal{F}$ by $\|f\|_2 = \sqrt{\langle f, f\rangle}$. The adjacency matrix A acts on $\mathcal{F}$ by

$$(Af)(v) = \sum_{w\sim v} f(w).$$

Since A is symmetric, there exists an orthonormal basis $\{f_1, \dots, f_n\}$ consisting of eigenfunctions. Let λ_i be the eigenvalue corresponding to f_i. If f, g are arbitrary functions, we can express f, g as combinations of eigenfunctions as $f = \sum_i \alpha_i f_i$, $g = \sum_i \beta_i f_i$. Then we have

$$\langle f, g\rangle = \langle \sum_i \alpha_i f_i, \sum_i \beta_j f_j\rangle = \sum_i, j\alpha_i\overline{\beta_j}\langle f_i, f_j = \sum_i \alpha_i\overline{\beta_i},$$

inparticular $\|f\|^2 = \sum_i |\alpha_i|^2$. Moreover

$$Af = \sum_i \alpha_i Af_i = \sum_i \lambda_i\alpha_i f_i.$$

Hence the decomposition of a function into eigenfunctions simplifies many computations.

3 Expanders

In reality graphs quite often occur as networks, where one can transport something along the edges, e.g. streets, railways, electricity lines, phone networks. In all these cases we want some sort of efficiency: On one hand, building many connections is expensive, on the other hand building too few connections results in networks unable to cope with the amount of traffic, or networks breaking down after little damage done by e.g. a storm. Of course, if all roads leading to one village are blocked, this village is isolated, but the damage is much smaller then if all main lines between two parts of the whole country are blocked. So what we really do not want is a network, which consists of different parts with few connections between them. In this case we speak of a bottleneck. Intuitively a bottleneck of a graph is a small set of edges, such that the graph obtained by deleting these edges is not only disconnected, but the largest connected component of the graph is much smaller then the whole graph.

Bottlenecks are either the result of poor planning or of natural contraints. The importance of the latter is obvious: Throughout history cities have been founded and battles fought over rivers, bridges, mountain passes ... . To prevent the former it is important to recognize bottlenecks. It follows from the physical picture introduced in the previous section that if $\mathcal{G}$ is a k-regular graph, and f is an eigenfunction to an eigenvalue very close to k, then there should be few connections between the set $\{v : f(v) > 0\}$ and the set $\{v : f(v) < 0\}$.

To make this intuition exact define for a set $A \subseteq V$ the boundary of A as $\partial A = \{v \in V : v \notin A, \exists a \in A : a \sim v\}$. Define the expansion constant of A as

$$\sup_{\substack{A \subseteq V \\ 1 \leq |A| \leq |V|/2}} \frac{|\partial A|}{|A|}.$$

A large expansion constant means that if a certain number n of lines break down, only a set of points of size proportional to n can become isolated. The relation between eigenvalues and expansion is given by the following.

Theorem 1. *Let $\mathcal{G}$ be a connected non-bipartite k-regular graph, λ_1 be the eigenvalue of largest absolute value different from λ. Then $\mathcal{G}$ has expansion constant $c \geq \frac{1}{2}\left(1 - \frac{\lambda_1}{k}\right)$.*

Proof. Since eigenfunctions are orthogonal, we may restrict A to the orthogonal complement of the eigenvector $\mathbf{1} = (1, 1, \ldots, 1)$. By the maximum principle we have

$$|\lambda_1| = \max_{f \neq 0 : \langle f, \mathbf{1} \rangle = 0} \frac{\langle Af, f \rangle}{\langle f, f \rangle}.$$

Let $A \subseteq V$ be a set satisfying $1 \leq |A| \leq \frac{|V|}{2}$. Then define a function

$$f_A(v) = \begin{cases} 1, & v \in A \\ -\frac{|A|}{|V|-|A|}, & v \notin A. \end{cases}$$

We have

$$\langle \mathbf{1}, f_A \rangle = \sum_{v \in A} 1 + \sum_{v \in V \setminus A} -\frac{|A|}{|V| - |A|} = 0,$$

thus $\frac{\langle Af, f \rangle}{\langle f, f \rangle} \leq \lambda_1$. We have

$$\langle Af, f \rangle = \sum_{v \sim w} f(v) f(w) = 2k|A| - (2 + \frac{2|V|}{|V| - |A|}) \# \{(v, w) : v \sim w, v \in A, w \notin A\} \geq 2k|A| - 4k|\partial A|,$$

while

$$\langle f, f \rangle = |A| + \frac{|A|^2}{|V| - |A|} \leq 2|A|.$$

Hence the relation between f and λ_1 implies $\frac{2k|A| - 4k|\partial A|}{2|A|} \leq \lambda_1$, and our claim follows. $\square$

Hence we are looking for graphs with λ_1 small. The best we could hope for is $\lambda_1 \approx \sqrt{k}$, in view of the following.

Theorem 2. *Let $\mathcal{G}$ be a k-regular graph with n vertices. Then the second largest absolute value of an eigenvalue of $\mathcal{G}$ is at least $\sqrt{\frac{kn-k^2}{n-1}}$.*

Proof. On one hand the trace of A^2 equals $2|E| = kn$, on the other hand the trace of A^2 equals the sum of the squares of the eigenvalues of A. Hence

$$kn = \sum_{\lambda \text{ Eigenvalue}} \lambda^2 \leq k^2 + (n-1)\lambda_1^2,$$

and our claim follows. $\square$

4 probabilistic constructions

It is difficult to think of an object which does not have any structure. In fact, as soon as we can describe something in an efficient way, this object must have a lot of structure. This implies that in many cases the mathematical objects we think of are not general, but quite special. To prove the existence of objects which cannot be described easily, one can resort to random constructions. One of the first random constructions is contained in the following.

Theorem 3. *Let $\mathcal{G}$ be a complete graph on $N = 2^{n/2}$ points. Then the edges of $\mathcal{G}$ can be coloured in red and blue in such a way, that there is no set A of n vertices, such that all edges between elements of A have the same colour.*

Proof. Colour the edges with equal probability red or blue. The probability that all edges among n vertices are red is $2^{-\frac{n(n-1)}{2}}$, hence the probability that all edges have the same colour is $2^{1-\frac{n(n-1)}{2}}$. If $\mathcal{G}$ has N vertices, then there are $\binom{N}{n}$ subsets of size n, hence the expected number of monochromatic subsets is

$$\binom{N}{n} 2^{1-\frac{n(n-1)}{2}} \leq \frac{2N^n}{n! 2^{\frac{n(n-1)}{2}}} \leq \frac{N^n}{2^{n^2/2}},$$

provided that $n \geq 4$. For $N < 2^{n/2}$ this is less than 1. If the expected number of monochromatic subgraphs is at most 1, some colouring contains no monochromatic subgraph, and our claim follows. For $n \leq 3$ the statement can be checked directly. □

Theorem 4. *For every large even integer n there exists a 3-regular graph with n vertices and expansion constant ≥ 0.09.*

Proof. We construct a 3-regular graph at random. To do so we partition the vertices at random into two sets of equal size, and connect one point from one set with a random point, which is not yet used from the other set. We repeat this process two more times. Let $A \subseteq B$ be sets with $|A| \leq |V|/2$ and $|B| \leq (1+c)|A|$. The probability that all edges starting in A have endpoint in B is $\leq \left(\frac{|B|}{|V|}\right)^{3|A|}$, hence the expected number of pairs A, B proving that $\mathcal{G}$ has expansion constant $< c$ is at most

$$\sum_{\nu=1}^{n/2} \binom{n}{\nu} \binom{n-\nu}{\lfloor c\nu \rfloor} \left(\frac{\lfloor (1+c)\nu \rfloor}{n}\right)^{3\nu}. \tag{1}$$

By Stirling's formula we have

$$\binom{n}{\nu} \binom{n-\nu}{\lfloor c\nu \rfloor} \leq \frac{n^{(1+c)\nu}}{(\nu/e)^\nu (c\nu/e)^{c\nu}} = \left(\frac{en}{\nu}\right)^{(1+c)\nu} c^{-c\nu}.$$

Hence,

$$\binom{n}{\nu} \binom{n-\nu}{\lfloor c\nu \rfloor} \left(\frac{\lfloor (1+c)\nu \rfloor}{n}\right)^{3\nu} \leq \left(e^{1+c} c^{-c} \left(\frac{\nu}{n}\right)^{2-c}\right)^{\nu}.$$

We have $\nu \leq n/2$, hence the expression in brackets on the right is $\leq \frac{e}{4}\left(\frac{2e}{c}\right)^c \leq 0.99$, provided that $c \leq 0.09$. This implies that if we restrict the summation in (1) to the range $500 \leq \nu \leq n/2$, the value of the sum is $\leq \sum_{\nu=500}^{\infty} 0.99^\nu \leq 0.658$. For smaller values of ν we can estimate $\frac{\nu}{n}$ better. If we assume that $n > 2500$, then we have in the missing range $\frac{\nu}{n} \leq \frac{1}{5}$, thus a summand is bounded above by $\left(\frac{e}{25}\left(\frac{5e}{c}\right)^c\right)^\nu$, for $c = 0.09$ this becomes $\leq 0.171^\nu$. The sum over the missing range is therefore ≤ 0.206, and we conclude that the whole sum is $\leq 0.864 < 1$. Hence the expected number of pairs of sets proving that $\mathcal{G}$ has expansion constant < 0.09 is less than 1, which implies that there exist some graph for which no such pair exists. But this graph has expansion ≥ 0.09, and our claim follows. □

Exercise. Try to improve the constant c! What happens with c when you increase k?

Sometimes an existence proof is fine, but sometimes one has to construct the object in question. Random constructions do not immediately give the latter. In our application we have seen that the probability of obtaining a graph with good expansion property is so large, that we can hope to get one by just producing many examples at random and checking. Note that already for very small graphs it is impossible to check all subsets of size $\leq n/2$, as there are 2^{n-1} of them. However, computing the eigenvalues of a large matrix is much easier, so even for constructing some real life network, we need Theorem 1.

5 Cayley graphs

Cayley graphs are examples of graphs which can easily be described, but using the correct parameters, they can look similar to random graphs.

Let G be a group, S a generating set, such that $1 \notin G$, and for all $s \in S$ we have $s^{-1} \in S$ as well. Then define a graph $\mathcal{G}$ by taking the elements of G as vertices, and the edges to be all sets $\{g, h\}$, such that $gh^{-1} \in S$. For example, if $G = C_n$ is cyclic of order n, and $S = \{\pm 1\}$, then $\mathcal{G}$ is a cycle of length n. If $G = C_n^2$, and $S = \{\pm\binom{1}{0}, \pm\binom{0}{1}\}$, then $\mathcal{G}$ can be visualized as an $n \times n$-grid, where each vertex in the upermost row is connected to the corresponding vertex in the lowermost row, and each vertex in the leftmost row is connected to the corresponding vertex in the rightmost row. These graphs are poor expanders. The reason is that there are many short relations among elements of S. More general in an abelian group we have $xy = yx$, that is, the Cayleygraph of an abelian group contains many cycles of length 4, while for a graph with good expansion properties we would rather have no short cycles at all.

If the generating set S is a conjugacy class, we can bound the eigenvalues of the Cayleygraph using representation theory.

Theorem 5. *Let G be a finite group, S a conjugacy class generating G, and assume that $s \in S$ implies $s^{-1} \in S$. Then the Cayleygraph $\mathcal{G}(G, S)$ has eigenvalues equal to $\frac{\chi(s)}{\chi(1)}$, where χ runs over all irreducible complex characters of G, and s is an arbitrary element of S. Moreover, the eigenvalue $\frac{\chi(s)}{\chi(1)}$ appears with multiplicity $\chi(1)^2$.*

The proof of this statement requires too much representation theory to be given here.

For practical purposes a generating conjugacy class will usually be too large to yield useful expanders, however, by choosing S more carefully one can in fact construct expanders of arbitrary size with bounded vertex degree. That bounded degree is sufficient, is the content of the following.

Lemma 1. *Let $\mathcal{G}$ be a c-expander in which all vertices have degree $\leq d$. Then there exists a $\frac{c}{d^2}$-expander $\mathcal{G}'$ with $|V| \leq |V'| \leq d|V|$, in which all vertices have degree 3.*

Proof. For each vertex v of $\mathcal{G}$ we order the outgoing edges cyclically in some way. We now define $\mathcal{G}'$ as follows. The vertex set V' consists of the set of directed vertices (v, v') of $\mathcal{G}$. Two vertices (v_1, v_1'), (v_2, v_2') are joined by an edge if either $(v_1, v_1') = (v_2', v_2)$, or $v_1 = v_2$, and the edges (v_1, v_1') and (v_2, v_2') are neighbours in the cyclical ordering chosen. Less formally if we view $\mathcal{G}$ as a set of cities connected by streets, then $\mathcal{G}'$ is obtained by building a ring road around each city, and calling the points where the old streets meet the ring road the new vertex set.

We claim that $\mathcal{G}'$ is a c-expander. Let A' be an arbitrary subset of the vertices of $\mathcal{G}'$. Define the map $\pi : V' \to V$ by mapping (v, v') to v. Let A be the set of vertices $v \in \mathcal{G}$, such that $\pi^{-1}(v) \subseteq A'$, and write $A' = \pi^{-1}(A)\dot{\cup}B$. If $v \in \pi(B)$, then at least one point of the set (v, v') is contained in A', and at least one point of this set is not contained in A'. Since $\pi^{-1}(v)$ is a connected, we obtain $\partial A' \cap \pi^{-1}(v) \neq \emptyset$, and therefore $|\partial A'| \geq |\pi(B)|$. Since each vertex of $\mathcal{G}$ has order $\leq d$, we conclude $|(\partial A') \cap B| \geq \frac{1}{d}|B|$.

If $v \in \partial A$, then either $v \in \pi(B)$, or some preimage of v is in the boundary of A, but not in B. We conclude that $|\partial A' \setminus B| \geq |\partial A| - |B|$. Hence we obtain

$$|\partial A'| = |\partial A' \setminus B| + |(\partial A') \cap B| \geq \max(0, c|A| - |B|) + \frac{|B|}{d} \geq \frac{c|A|}{d} \geq \frac{c|A'|}{d^2}.$$

Hence our claim follows. $\square$

If one is looking for exapnders with vertex order d different from 3, one can use the same method, but replace each vertex not by a cycle but by an expander with vertex order $d-1$. In this way the poor expansion property of a long road can be circumvented at the expense of getting higher vertex degrees.

Exercise. Try different recursive constructions. How far can you improve the dependence on d in the Lemma, if you want to get a 4-regular graph, a 5-regular graph, or in general an r-regular graph?