Graph Theory: CMSC 27530/37530 Lecture 4

Lecture by László Babai
Notes by Geoffrey West

April 11, 2019*

Definition 4.1. For two vertices $x, y \in V$, the **distance** from x to y, denoted $\text{dist}(x, y)$, is the length of the shortest $x...y$ path. If there is no $x...y$ path, we say $\text{dist}(x, y) = \infty$.

Note that the distance defines a metric on the set of vertices:

(i) $\text{dist}(x, y) = \text{dist}(y, x)$

(ii) $\text{dist}(x, y) \geq 0$, with equality if and only if $x = y$.

(iii) $\text{dist}(x, z) \leq \text{dist}(x, y) + \text{dist}(y, z)$ (Triangle Inequality).

DO 4.2. Show that $\text{dist}(,)$ satisfies the triangle inequality.

Definition 4.3. The **diameter** of a graph G is

$$\text{diam}(G) = \max_{x,y \in V}\{\text{dist}(x, y)\}.$$

DO 4.4. Show the following.

(a) $(\forall G)(G \text{ or } \overline{G} \text{ is connected})$.

(b) $\min\{\text{diam}(G), \text{diam}(\overline{G})\} \leq 3$.

(c) If $\text{diam}(G) \geq 4$, then $\text{diam}(\overline{G}) \leq 2$.

DO 4.5. Show that $(c) \Rightarrow (b) \Rightarrow (a)$.

Recall that we have previously seen $\max_n\{m : G \text{ is disconnected}\} = \binom{n-1}{2}$. How can we prove this? The graph $K_{n-1} \cup K_1$ achieves this number. To show the upper bound, we need to show that if G is disconnected then $m_G \leq \binom{n-1}{2}$, i. e., $m_{\overline{G}} \geq \binom{n}{2} - \binom{n-1}{2} = n - 1$. But this follows from the fact that $\overline{G}$ must be connected in view of exercise DO 4.4.

Recall the Mantel–Turán: a triangle-free graph has at most $\lfloor \frac{n^2}{4} \rfloor$ edges. For a bipartite graph,

$$\max\{m : G \text{ is bipartite with } n \text{ vertices}\} = \max_{0 \leq k \leq n} k(n-k) = \lfloor \frac{n^2}{4} \rfloor.$$

*Posted April 11 at 7pm. Revised April 12, 8am.

This is the same as the result given in Mantel–Turán. Note that every bipartite graph is K_3-free, so

$$\{K_3\text{-free graphs}\} \supset \{\text{bipartite graphs}\}.$$

As a result,

$$\max\{m : K_3\text{-free with } n \text{ vertices}\} \geq \max\{m : \text{bipartite with } n \text{ vertices}\}.$$

Mantel–Turán shows that, somewhat surprisingly, the two sides are actually equal.

DO 4.6. The maximum number of edges (given n) for a K_3-free graph occurs only if the graph is bipartite, i.e. $K_{\lfloor \frac{n}{2} \rfloor, \lceil \frac{n}{2} \rceil}$.

Similarly, we observe that $\max_n\{m : K_{r+1}\text{-free}\} \geq \max_n\{m : r\text{-colorable}\}$. The following theorem shows equality.

Theorem 4.7 (Turán's Theorem)**.** $\max_n\{m : K_{r+1}\text{-}\textit{free}\} = \max_n\{m : r\text{-}\textit{colorable}\}$.

Definition 4.8. The **complete r-partite graph** $K_{n_1,\dots,n_r}$ has $n = n_1 + \cdots + n_r$ vertices ($n_i \geq 1$) divided into r parts where the i-th part has n_i vertices and two vertices are adjacent precicely if they don't belong to the same part.

Note that for $r = 2$ we get the definition of the complete bipartite graphs.

DO 4.9. If G is a complete r-partite graph then G is r-colorable. Moreover, G is a *maximal* r-colorable graph in the sense that it ceases to be r-colorable if we add any edge.

DO 4.10. Find $\max_n\{m : r\text{-colorable}\}$. Show that this occurs for the complete r-partite graph with almost equal parts: each part contains either $\lfloor \frac{n}{r} \rfloor$ or $\lceil \frac{n}{r} \rceil$ vertices.

DO 4.11. Prove Turán's theorem by induction on n. First saturate the graph with respect to the property of not containing a K_{r+1} : keep adding edges as long as you can without creating a K_{r+1}. Note that the saturated graph contains K_r. Pick a subgraph $K_r \subseteq G$ and induct in a manner similar to our inductive proof of Mantel–Turán. The base cases are $n = 0, 1, ..., r-1$.

Proposition 4.12. $\max_n\{m : r\text{-}\textit{colorable}\} \leq (1 - \frac{1}{r})\frac{n^2}{2}$.

HW 4.13. (4 points) Prove this Proposition. The proof should be no more than three lines. (Four if you are verbose.)

Combining this with Theorem 4.7 we obtain the following.

Corollary 4.14 (Turán)**.** *If G is K_{r+1}-free, then* $m \leq \left(1 - \frac{1}{r}\right)\frac{n^2}{2}$.

Next we derive an interesting consequence of Turán's Theorem regarding the **independence number** $\alpha(G)$. First we make a simple observation. Δ is the maximum degree.

Observation 4.15 (Naive bound on the independence number)**.** $\alpha(G) \geq \dfrac{n}{1+\Delta}$.

Proof 1. Based on previous exercises, $\alpha(G) \geq \frac{n}{\chi(G)} \geq \frac{n}{\Delta + 1}$. □

Proof 2. Use a greedy independent set algorithm. Let $V = \{v_1, ..., v_n\}$.
  initialize $I := \varnothing$.
  for $i = 1$ **to** n:
    if v_i has no neighbor in I **then** $I \leftarrow I \cup \{v_i\}$
  end(for)
  return I. □

DO 4.16. Show that the greedy algorithm returns an independent set of size $\geq \frac{n}{1 + \Delta}$.

Next we strengthen this result by replacing the maximum degree by the average degree $(\sum d_i)/n = 2m/n$. It turns out that this stronger lower bound on $\alpha(G)$ is an immediate consequence of Turán's Theorem.

Theorem 4.17 (Turán's bound on the independence number)**.** $\alpha(G) \geq \frac{n}{1 + \frac{\sum d_i}{n}}$.

HW 4.18. (4 points) Derive this inequality from Corollary 4.14. Your proof should be no more than 3 lines.

Next we state a remarkable further strenthening of the bound.

Theorem 4.19 (Victor Keh-Wei Wei and Yair Caro, independently, cca. 1980)**.**

$$\alpha(G) \geq \sum_{i=1}^{n} \frac{1}{1 + d_i}.$$

Note: Wei's full name was tracked down by your classmates Jeremiah Milbauer, Shashank Srivastava, and Can Liu from websites at the City University of Hong Kong and a biographical sketch that appeared in a 1995 article.

We shall see that the Wei–Caro bound is stronger than Turán's. Its proof is a gem.

Definition 4.20. If $x_1, ..., x_n > 0$, their **geometric mean** is

$$G(x_1, ..., x_n) = \left(\prod_{i=1}^{n} x_i\right)^{1/n}.$$

DO 4.21. Show that $G \leq A$, where A is the arithmetic mean. Equality holds if an only if all the x_i are equal.

Hint for a high-school solution: first prove it for $n = 2^k$ by induction on k. Then prove that if we know the inequality for some value of n then the inequality follows for all smaller values of n.

A more advanced proof uses Jensen's inequality and the concavity of the logarithm function.

If you have difficulty executing these hints, find a couple of proofs online.

Definition 4.22. If $x_1, ..., x_n > 0$, the **harmonic mean** is

$$H(x_1, ..., x_n) = \frac{1}{A(\frac{1}{x_i}, ..., \frac{1}{x_n})} = \frac{n}{\sum \frac{1}{x_i}}.$$

Theorem 4.23. *$G \geq H$. Equality holds if and only if $x_1 = \ldots = x_n$.*

HW 4.24. (3 points) Prove this inequality, in one line, using $G \leq A$.

Corollary 4.25. *$A \geq H$.*

HW 4.26. (3 points) Prove that the Wei–Caro bound is always at least as strong as Turán's, i. e., prove that

$$\sum_{i=1}^{n} \frac{1}{1+d_i} \geq \frac{n}{1 + \frac{\sum_{i=1}^{n} d_i}{n}}.$$

Example 4.27. To illustrate how much stronger Wei–Caro can be than Turán, consider a graph with two connected components of order $n/2$ each, one of them 3-regular, the other a clique. Then the average degree is greater than $n/4$. Therefore the Turán lower bound is not greater than 4. Wei–Caro, on the other hand, shows that $\alpha \geq \frac{n}{2}\frac{1}{1+3} + \frac{n}{2}\frac{1}{n/2} > \frac{n}{8}$.

To prove Wei–Caro, we need to take an excursion into

random variables over **finite probability spaces.**

Consider shuffling a deck of 52 cards. There are 52! possible outcomes of this experiment. If we flip n coins, we get a binary sequence (Heads/Tails) of length n. There are 2^n possible outcomes of this experiment. We can construct a "random graph" by fixing a set of n vertices and deciding adjacency by flipping a (possibly biased) coin for each pair of vertices. This experiment has $2^{\binom{n}{2}}$ possible outcomes.

We refer to the set of possible outcomes of each experiment described as the *sample space* of the given experiment. This illustrates the first component of the abstract concept of a *probability space*. We refer to each outcome of the experiment (each element of the sample space) as an "elementary event." We intuitively associate a "probability" with each elementary event. This illustrates the second component of the concept of a probability space. After this motivation, we give the abstract definition in terms of two basic mathematical primitives, sets and functions. "Experiment" is not a mathematical concept and it is not mentioned in the definition – the examples involving "experiments" only serve to help our intuition.

Definition 4.28. A **probability distribution** over a non-empty finite set Ω is a function $P : \Omega \to \mathbb{R}$ satisfying

(i) $(\forall x \in \Omega)(P(x) \geq 0)$

(ii) $\sum_{x \in \Omega} P(x) = 1$.

Definition 4.29. A **finite probability space** is a pair (Ω, P) where Ω is a non-empty finite set, called the **sample space**, and P is a probability distribution over Ω.

Definition 4.30 (Probability of event)**.** An **event** is a subset $A \subseteq \Omega$ of the sample space. The probability of an event A is defined as

$$P(A) = \sum_{x \in A} P(x).$$

DO 4.31. $P(\varnothing) = 0$ and $P(\Omega) = 1$.

Definition 4.32. A **trivial event** is an event with probability 0 or 1.

DO 4.33 (Complement)**.** For an event A, we have $P(\overline{A}) = 1 - P(A)$, where $\overline{A} = \Omega \setminus A$.

DO 4.34. If $A \cap B = \varnothing$, then $P(A \cup B) = P(A) + P(B)$.

DO 4.35 (Modular equation)**.** If $A, B \subseteq \Omega$ are events then

$$P(A \cup B) + P(A \cap B) = P(A) + P(B).$$

Definition 4.36 (Union bound)**.** If $A_1, \ldots, A_k \subseteq \Omega$ are events then

$$P\left(\bigcup_{i=1}^{k} A_i\right) \leq \sum_{i=1}^{k} P(A_i).$$

Definition 4.37. P is the **uniform distribution** if $(\forall x \in \Omega)\left(P(x) = \frac{1}{|\Omega|}\right)$.

DO 4.38. If P is uniform then the probability of event A is $\frac{|A|}{|\Omega|}$ ("number of good cases divided by the number of all cases") – the naive notion of probability.

Convention 4.39 (Default distribution: uniform)**.** To specify a probability space, we need to state the sample space and the probability distribution. If we omit the probability distribution, then the probability distribution is assumed to be uniform.

Definition 4.40. A **random variable** is a function $X : \Omega \to \mathbb{R}$.

Example 4.41. The number of heads in a random sequence of n coin flips is a random variable. (The sample space is the set of possible coin-flip sequences.)

Definition 4.42. For a random variable X, the **expected value** of X is the quantity

$$E(X) = \sum_{x \in \Omega} X(x) \cdot P(x). \tag{1}$$

Note that this is a weighted average of the values of X.

DO 4.43. Show that

$$\min X \leq E(X) \leq \max X.$$

DO 4.44. If the distribution P is uniform then the expected value is simply the average of the values taken by X:

$$E(X) = \frac{1}{|\Omega|} \sum_{x \in \Omega} X(x).$$

The most important fact about the expected value is that it is linear:

(i) $E(X + Y) = E(X) + E(Y)$

(ii) $E(c \cdot X) = c \cdot E(X)$

where X and Y are random variables over the same probability space $)\Omega, P)$ and c is a scalar (real number).

DO 4.45 (Linearity of Expectation)**.** Show that

$$E\left(\sum c_i X_i\right) = \sum c_i E(X_i).$$

(This is equivalent to the two preceding statements.)

The sample space is often much larger than the range of values of a random variable. (The random variable may take the same value on large subsets of the sample space.) Grouping together those terms in Eq. (1), we obtain the expression in the next exercise; this has only as many non-zero terms as the size of the range of X.
Notation: The expression "$X = y$" denotes the event $\{x \in \Omega \mid X(x) = y\}$.

DO 4.46. Prove: $E(X) = \sum_{y \in \mathbb{R}} y \cdot P(X = y)$.

Definition 4.47. An **indicator variable** is a random variable that takes values 0 and 1 only.

There is a 1-1 correspondence between indicator variables and events. Let $A \subseteq \Omega$ be an event. We define the corresponding indicator variable ϑ_A as follows. For $x \in \Omega$ we set

$$\vartheta_A(x) = \begin{cases} 1 & x \in A \\ 0 & x \notin A \end{cases}$$

DO 4.48. Find the inverse of this correspondence: show that if X is an indicator variable then letting $A = X^{-1}(1)$ we have $X = \vartheta_A$.

Remark 4.49. Events have probabilities; they don't have expected values. Random variables have expected values; they don't have probabilities.

Next we illustrate the power of the linearity of expectation on a problem.
The hat-check problem, first variant.

On entering a club, n patrons check their hats at a counter. On leaving, the hat-check clerk returns a random hat to each patron. What is the expected number of patrons who get their own hat?

Let us formalize this question.

Definition 4.50. A **permutation** of a set S is a bijection $f : S \to S$. A **fixed point** of f is a value $x \in S$ such that $f(x) = x$.

So the question is, what is the expected number of fixed points of a random permutation.

To further formalize this question, we consider the uniform probability space of which the sample space Ω is the set of $n!$ permutations of a set S of n elements. For $f \in \Omega$, let $X(f)$ denote the number of fixed points of f. So X is a random variable and the question is to find $E(X)$.

To do so, we use indicator variables. Let $S = \{s_1, \ldots, s_n\}$. Let Y_i be the indicator of the event $f(s_i) = s_i$ (s_i is a fixed point). Then

$$X = Y_1 + \ldots + Y_n \tag{2}$$

DO 4.51. Verify Eq. (2). What it means is that for every $f \in \Omega$ we have $X(f) = \sum_{i=1}^{n} Y_i(f)$.

From Eq. (2) it follows, by the linearity of expectation, that

$$E(X) = \sum_{i=1}^{n} E(Y_i). \tag{3}$$

Now $E(Y_i) = P(f(s_i) = s_i)$ (the probability of the event indicated by Y_i). This probability is $1/n$ by symmetry: $f(s_i)$ could be any member of S, with equal probability, so the probability that s_i is fixed by f is $1/n$.

Summarizing,

$$E(X) = \sum_{i=1}^{n} E(Y_i) = \sum_{i=1}^{n} \frac{1}{n} = 1.$$

So the expected number of hats returned to their owner is 1.

BONUS 4.52. (6 points) Prove Wei–Caro in this manner.
Hint. Given a graph G, first you need to construct a probability space. State the size of the sample space. Then you need to construct a random variable X such that

- $X \leq \alpha(G)$
- $E(X) = \sum \frac{1}{1 + d_i}$

Use indicator variables to compute the expected value of your random variable X. This hint should suffice. DO NOT spoil your fun by looking up the solution online.