BRICS LS-95-4 Breslauer & Dubhashi: Combinatorics for Computer Scientists

BRICS

Basic Research in Computer Science

Combinatorics for Computer Scientists

Dany Breslauer
Devdatt P. Dubhashi

BRICS Lecture Series **LS-95-4**

ISSN 1395-2048 **August 1995**

See back inner page for a list of recent publications in the BRICS
Lecture Series. Copies may be obtained by contacting:

BRICS
Department of Computer Science
University of Aarhus
Ny Munkegade, building 540
DK - 8000 Aarhus C
Denmark

Telephone: +45 8942 3360
Telefax: +45 8942 3255
Internet: BRICS@daimi.aau.dk

BRICS publications are in general accessible through WWW and
anonymous FTP:

`http://www.brics.aau.dk/BRICS/`
`ftp ftp.brics.aau.dk (cd pub/BRICS)`

Combinatorics for Computer Scientists

Dany Breslauer
Devdatt P. Dubhashi

Dany Breslauer
Devdatt P. Dubhashi
BRICS[1]
Department of Computer Science
University of Aarhus
Ny Munkegade
DK-8000 Aarhus C, Denmark

[1] **B**asic **R**esearch **I**n **C**omputer **S**cience, Centre of the Danish National Research Foundation.

Preface

These are informal lecture notes from the course *Combinatorics for Computer Scientists* that was offered at Aarhus University in Spring 1995. As a textbook that covers many of the topics we were interested in, we recommeneded:

- J.H. van Lint and R.M. Wilson, *A Course in Combinatorics*, Cambridge University Press, 1992,

but we extensively supplemented it with class notes. The notes, for most part, were taken by the students participating in the class: Peer Duschner, Claus Gregersen, Thomas Hildebrandt, Lone Overgaard, Theis Rauhe, Søren Skyum and Peter Unold. We have then edited the notes, sometimes reordering parts and organizing them in a more uniform format.

The material covered in the course can be grouped into roughly three separate parts corresponding to the three parts in which these notes are organised.

- *Enumerative Combinatorics*: The principal topics covered were the Inclusion–Exclusion formula, Möbius Inversion and Generating Functions. Aside from the textbook, the principal references we consulted for this part were:
  - L. Lovász, *Combinatorial Problems and Exercises* (2nd edition), North Holland Publishing Company and Akadémiai Kiodó, 1993.
  - D.E. Knuth R. L. Graham and O. Patashnik, *Concrete Mathematics* (2nd edition), Addison Wesley, 1994.
  - H. Wilf, *Generatingfunctionology* (2nd edition), Academic Press, 1994, .
  - R.P. Stanley, *Enumerative Combinatorics*, Part I, Wadsworth and Brooks/Cole, 1986.

  This part was taught mainly by DD
- *Graph Theory*: This consisted of a fairly standard set of topics in Graph Theory: The main reference we used for this part was the following text
  - J.A. Bondy and U.S.R. Murty, *Graph Theory with Applications*, Macmillan Press, 1977.

  This part was taught mainly by DB.
- *Linear Algebra Methods*: This was a somewhat novel part of the course in which we aimed to show how some fairly elementary dimension arguments from Linear Algebra could be taken quite far to yield non–trivial results in combinatorics. The method was illustrated via a series of results related to *extremal set theory*. A nice application that tied in with the graph theory part was a constructive proof of a Ramsey lower bound. The material was mainly based on the following set of notes which is to appear as a textbook in the near future:
  - L. Babai and P. Frankl, *Linear Algebra Methods in Combinatorics*, University of Chicago, Department of Computer Science.

Also for the part on matroids, we were influenced by the treatment in Chapter 3 of

– D. Kozen, *The Design and Analysis of Algorithms*, Springer–Verlag, 1992.

This part was taught by DD.

The notes for the corresponding parts reflect our individual stylistic differences.

Dany Breslauer and Devdatt P. Dubhashi
Aarhus, May 1995

Contents

Part I

Enumeration

Chapter 1

Inclusion–Exclusion

In the proper spirit of combinatorics, let us begin with one of the simplest counting problems possible: of the 100 members of DAIMI, 70 like Tuborg and 65 Carlsberg. How many like both? [1] The answer, 35, is obtained by appyling the formula:

$$|A \cup B| = |A| + |B| - |A \cap B|, \tag{1.1}$$

valid for two arbitrary finite subsets of a universal set, U.

Let us demonstrate why this holds; it will give us a chance to introduce some handy notation. First, for any subset $X \subseteq U$, we have,

$$\begin{aligned} |X| &= \sum_{x \in X} 1 \\ &= \sum_{x \in U} [x \in X] \\ &= \sum_{x} [x \in X]. \end{aligned}$$

Notice that a summation index without conditions ranges over all possible values. Also, we have commandeered the *Iverson* notation from APL,

$$[P] := \begin{cases} 1, & \text{if } P \text{ is true;} \\ 0, & \text{otherwise.} \end{cases}$$

for an arbitrary boolean property, P. Then, (1.1) is equivalent to the assertion that

$$\sum_x [x \in A] + [x \in B] - [x \in A \cup B] - [x \in A \cap B] = 0.$$

This sum is indeed zero for the trivial reason, that each summand,
$f(x) := [x \in A] + [x \in B] - [x \in A \cup B] - [x \in A \cap B]$ is, as the following *truth table* confirms:

$[x \in A]$	$[x \in B]$	$[x \in A \cup B]$	$[x \in A \cap B]$	$f(x)$
0	0	0	0	0
0	1	1	0	0
1	0	1	0	0
1	1	1	1	0

[1] Silly objections that some people might not drink at all can be dismissed in Denmark!

What happens if we add a third brand of Öl? The formula is now,

$$|A \cup B \cup C| = |A| + |B| + |C| - |B \cap C| - |C \cap A| - |A \cap B| + |A \cap B \cap C|.$$

In general, with $n \geq 1$ brands, that is, for $A_1, \ldots A_n \subseteq U$, the formula is:

$$|A_1 \cup \cdots \cup A_n| = \sum_i |A_i| - \sum_{i<j} |A_i \cap A_j| + \cdots + (-1)^{n-1}|A_1 \cap \cdots \cap A_n|.$$

More compactly, we can write this formula as follows[2]:

$$|\bigcup_{i\in[n]} A_i| = \sum_{\emptyset \neq I \subseteq [n]} (-1)^{|I|-1} |\bigcap_{i \in I} A_i|. \tag{1.2}$$

This formula is called the *inclusion–exclusion formula*, and is one of the most basic principles in combinatorics. We shall postpone the proof of equation (1.2) so that we can illustrate how it is used.

There is another form of the equation (1.2) that is sometimes useful. Suppose that we are in a *symmetric* situation where $|\bigcap_{i\in I} A_i| = |\bigcap_{j \in J} A_j|$ whenever $|I| = |J|$ for $I, J \subseteq [n]$. Thus, if $|I| = k$, then $|\bigcap_{i\in I} A_i|$ depends only on k. Denote this value by a_k. We use the notation $\binom{[n]}{k}$ to denote the set of all subsets of $[n]$ of size k. We have,

$$\begin{aligned}
\sum_{\emptyset \neq I \subseteq [n]} (-1)^{|I|-1} |\bigcap_{i\in I} A_i| &= \sum_{1 \leq k \leq n} \sum_{I \in \binom{[n]}{k}} (-1)^{|I|-1} |\bigcap_{i \in I} A_i| \\
&= \sum_{1 \leq k \leq n} \sum_{I \in \binom{[n]}{k}} (-1)^{k-1} a_k \\
&= \sum_{1 \leq k \leq n} (-1)^{k-1} a_k \sum_{I \in \binom{[n]}{k}} 1 \\
&= \sum_{1 \leq k \leq n} (-1)^{k-1} a_k \binom{n}{k}.
\end{aligned}$$

Hence, in the symmetric case, the inclusion–exclusion formula reads:

$$|\bigcup_{i\in[n]} A_i| = \sum_{1 \leq k \leq n} (-1)^{k-1} a_k \binom{n}{k}. \tag{1.3}$$

Example 1 A hundred Danish fans go to see the European football championships. A Brian and Michael Laudrup combination sets up a goal, and the fans go wild! They fling their hats into the air. A gust of wind catches the hats and returns them in a random order to their owners. One of the fans who is a DAIMI professor ponders on the following perverse question: how many ways are there to return the hats in such a way that nobody gets his/her own hat back?

Less frivolously, a *permutation* on $n \geq 1$ elements is a one–one and onto mapping $\pi : [n] \to [n]$. We denote the set of all permutations on n elements by $\mathcal{S}_n$. There are a total of $n!$ permutations in $\mathcal{S}_n$. A *derangement* is a permutation π, such that $\pi(i) \neq i$ for every $i \in [n]$. The professor's question is: how many derangements, d_n are there in $\mathcal{S}_n$?

[2] We shall use the notation $[n] := \{1, \ldots, n\}$ consistently.

We use the inclusion–exclusion formula applied to the universe $\mathcal{S}_n$. It is crucial to choose the sets A_i judiciously – so that computing $|\bigcap_{i \in I} A_i|$ is easy for every $I \subseteq [n]$. To this end, we choose, for each $i \in [n]$,

$$A_i := \{\pi \in \mathcal{S}_n \mid \pi(i) = i\}.$$

Then,

$$\bigcap_{i \in I} A_i = \{\pi \in \mathcal{S}_n \mid \pi(i) = i \quad \text{for } i \in I\}.$$

Hence,

$$|\bigcap_{i \in I} A_i| = (n - |I|)!$$

and we are in a symmetric case. Applying equation (1.3), we get:

$$\begin{aligned} |\bigcup_i A_i| &= \sum_{1 \le k \le n} (-1)^{k-1} (n-k)! \binom{n}{k} \\ &= n! \cdot \sum_{1 \le k \le n} \frac{(-1)^{k-1}}{k!}. \end{aligned}$$

Subtracting this from the total number of permutations, the number of derangements is

$$d_n = n! \sum_{0 \le k \le n} \frac{(-1)^k}{k!}. \tag{1.4}$$

It is useful to approximate this by extending the summation over all k,

$$\begin{aligned} d_n &\approx n! \sum_{k \ge 0} \frac{(-1)^k}{k!} \\ &= n!/e \end{aligned} \tag{1.5}$$

Exercise 2 *Check this approximation to show that in fact d_n is* exactly *equal to the integer nearest to $n!/e$, i.e. for $n \ge 1$,*

$$d_n = \lfloor n!/e + 1/2 \rfloor.$$

We might ask the original question as follows: what is the *probability* that nobody gets his/her own hat back? Such a question can be asked in the setting of a *probability space* which is a set Ω, called the *sample space* and a *probability distribution* which is a function $\Pr : \Omega \to [0, 1]$ such that

$$\sum_{\omega \in \Omega} \Pr(\omega) = 1.$$

An *event* is a subset $A \subseteq \Omega$, and the probability of the event is

$$\Pr(A) := \sum_{\omega \in A} \Pr(\omega) = \sum_{\omega} \Pr(\omega) \cdot [\omega \in A].$$

In many situations, the space Ω is finite and each point $\omega \in \Omega$ is democratically assigned the same value, $\Pr(\omega) := 1/|\Omega|$. In this case, the probability of an event $A \subseteq \Omega$ is given by

$$\Pr(A) = |A|/|\Omega|.$$

This is indeed the case in our situation. Our sample space is $\Omega := \mathcal{S}_n$ and each point, a permutation π is assigned the same probability $1/n!$. The event of interest corresponds to the subset of derangements, and so the probability of this event is, by equation (1.5), very nearly $1/e \approx 0.367$, *independent* of n.

Exercise 3 *Let A, B be events in a probability space. Show that*

$$\Pr(A \cup B) = \Pr(A) + \Pr(B) - \Pr(A \cap B).$$

How does this generalise?

Example 4 Our second example is from Number Theory. For a positive integer n, the *Euler totient function*, $\phi(n)$ is defined as the number of integers between 1 and n that are *relatively prime* to n:

$$\phi(n) := |\{k \in [n] \mid (n,k) = 1\}|.$$

Can we give a formula to compute ϕ?

The answer is easy in the case when n is a *prime* number, p:

$$\phi(p) = p - 1. \tag{1.6}$$

If n is a power of a prime, $n := p^\alpha$ where α is a positive integer, then an integer $k \in [n]$ is not relatively prime to n exactly when p is a factor of k. The other factor can be anything from 1 upto $p^{\alpha-1}$, so

$$\begin{aligned}\phi(p^\alpha) &= p^\alpha - p^{\alpha-1} \\ &= p^\alpha(1 - \frac{1}{p}). \end{aligned} \tag{1.7}$$

Note that this agrees with equation (1.6) when $\alpha = 1$.

For general n, we are naturally led to consider the unique *prime factorisation*

$$n := p_1^{\alpha_1} \cdots p_r^{\alpha_r},$$

where $p_1, \ldots p_r$ are distinct primes and $\alpha_1, \ldots \alpha_r$ are non–negative integers, for some $r \geq 1$. Knowing equation (1.7), we are strongly tempted to rush in (where angels fear to tread!) and compute:

$$\begin{aligned}\phi(n) &= \phi(p_1^{\alpha_1} \cdots p_r^{\alpha_r}) \\ &= \phi(p_1^{\alpha_1}) \cdots \phi(p_r^{\alpha_r}) \\ &= p_1^{\alpha_1} \cdots p_r^{\alpha_r} \cdot (1 - \frac{1}{p_1}) \cdots (1 - \frac{1}{p_r}) \quad \text{using eq. (1.7)} \\ &= n(1 - \frac{1}{p_1}) \cdots (1 - \frac{1}{p_r}) \\ &= n \prod_i (1 - \frac{1}{p_i}). \end{aligned}$$

That is,

$$\phi(n) = n \prod_i (1 - \frac{1}{p_i}). \tag{1.8}$$

This is indeed true – provided we verify that ϕ has the following *multiplicative* property:

$$\phi(n_1 \cdot n_2) = \phi(n_1) \cdot \phi(n_2), \quad \text{if } (n_1, n_2) = 1. \tag{1.9}$$

Although there is a direct way to argue that this property holds (see the next exercise), we shall prove it in a rather more interesting manner.

Exercise 5 *Argue directly that the Euler totient has the multiplicative property (1.9).*

Exercise 6 *Consider the setting of the general inclusion–exclusion formula, (1.2). For each $k \in [n]$, set*

$$S_k := \sum_{I \in \binom{[n]}{k}} |\bigcap_{i \in I} A_i|.$$

Prove the following Bonferroni inequalities *for each $p \geq 1$:*

$$\sum_{k \in [2p]} (-1)^{k-1} S_k \leq |\bigcup_{i \in [n]} A_i| \leq \sum_{k \in [2p-1]} (-1)^{k-1} S_k.$$

Exercise 7 *The lecturers share a house that has a very simple ground plan: the house is laid out on an east–west axis, with two bedrooms at either end and a shared kitchen and bath in between. Naturally, there are doors that connect each bedroom to both the kitchen and the bath. What is the minimum number of colours required to paint the rooms of the house so that every two rooms connected by a door are painted with different colours?*

Exercise 8 *By counting the number of onto mappings from a set of n elements to one of k elements, deduce the following combinatorial identity of Euler valid for all integers n, k:*

$$\sum_{0 \leq i \leq k} (-1)^i \binom{k}{i} (k-i)^n = \begin{cases} n!, & \text{if } k = n; \\ 0, & \text{if } k > n. \end{cases}$$

Chapter 2

Inclusion–Exclusion II

Last time, we were left with the task of showing that the Euler totient function ϕ is multiplicative. Recall that a function $f : \mathsf{N} \to \mathsf{N}$ is multiplicative if it satisfies (1.9). (Also to avoid triviality, insist that $f(1) = 1$: otherwise f will be identically zero!).

We'll prove this via an interesting detour. Consider $n = 12$ and write out the following fractions:

$$\frac{1}{12}, \frac{2}{12}, \frac{3}{12}, \frac{4}{12}, \frac{5}{12}, \frac{6}{12}, \frac{7}{12}, \frac{8}{12}, \frac{9}{12}, \frac{10}{12}, \frac{11}{12}, \frac{12}{12}.$$

Reduce the fractions and group them by denominators:

$$\frac{1}{1}, \frac{1}{2}, \frac{1}{3}, \frac{2}{3}, \frac{1}{4}, \frac{3}{4}, \frac{1}{6}, \frac{5}{6}, \frac{1}{12}, \frac{5}{12}, \frac{7}{12}, \frac{11}{12}.$$

Observe that by counting these fractions in two different ways, we see that: $12 = \phi(1)+\phi(2)+\phi(3)+\phi(4)+\phi(6)+\phi(12)$. The same argument shows that in general, ϕ has the following very important property:

$$n = \sum_{d|n} \phi(d). \tag{2.1}$$

We shall see a general technique to *invert* this equation to determine ϕ later.

Clearly the identity function $I(n) := n$ on the left side of (2.1) is multiplicative. Suppose we have two functions $f, g : \mathsf{N} \to \mathsf{N}$ that are related in the same way as in (2.1):

$$g(n) = \sum_{d|n} f(d), \tag{2.2}$$

(and furthermore $g(1) = 1 = f(1)$.) It is an interesting fact that in such a situation, whenever g is multiplicative, then so is f! That is, if g satisfies (1.9), then so does f. We shall prove this by induction on $n_1 \cdot n_2$. If $n_1 \cdot n_2 = 1$, then there is nothing to prove since $f(1) = 1$. By induction, suppose f satisfies (1.9) whenever $n_1 \cdot n_2 < n$ for some $n > 1$ and let us compute, for a pair of relatively prime positive integers n_1, n_2 with $n_1 \cdot n_2$,

$$\begin{aligned} g(n_1 n_2) &= \sum_{d|n_1 n_2} f(d) \\ &= \sum_{d_1|n_1} \sum_{d_2|n_2} f(d_1 d_2), \quad \text{since } (n_1, n_2) = 1. \\ &= \sum_{d_1|n_1} \sum_{d_2|n_2} f(d_1) \cdot f(d_2) - f(n_1) f(n_2) + f(n_1 n_2) \quad \text{by induction.} \end{aligned}$$

$$
\begin{aligned}
&= \sum_{d_1|n_1} f(d_1) \sum_{d_2|n_2} f(d_2) - f(n_1)f(n_2) + f(n_1n_2), \\
&= g(n_1)g(n_2) - f(n_1)f(n_2) + f(n_1n_2)
\end{aligned}
$$

But – and this is the clincher! – $g(n_1n_2) = g(n_1)g(n_2)$ since g is multiplicative. Hence $f(n_1n_2) = f(n_1)f(n_2)$ as well, and this completes the inductive step. This general fact implies, via (2.1) that ϕ is multiplicative.

Exercise 9 *Prove the following stronger converse to the general assertion in the previous paragraph: Suppose $f, g, h : \mathsf{N} \to$ integers are related by*

$$g(n) = \sum_{d|n} f(d)h(\frac{n}{d}),$$

and that f and h are multiplicative. Then so is g. (Note that with $h(n) := 1$, this gives the converse of the statement above.)

We shall now proceed to more general and abstract settings. A *partially ordered set* or *poset* for short, is a structure $\mathsf{P} := (P, \leq)$ where P is a set and $\leq$ is a binary relation on P called a *partial order*. The relation $\leq$ is:

1. *refelxive*, that is, $A \leq A$ for each $A \in P$,
2. *anti–symmetric*, that is, $A \leq B$ and $B \leq A$ implies $A = B$ for all $A, B \in P$, and
3. *transitive*, that is, $A \leq B$ and $B \leq C$ implies $A \leq C$ for all $A, B, C \in P$.

A special kind of poset is a *lattice* which is a structure $\mathsf{L} := (L, \vee, \wedge, \perp, \top)$. The operations $\vee$ and $\wedge$ are binary functions on L and are called the *join* and *meet* respectively. For $A, B \in L$ the element $A \vee B \in L$ satisfies:

- $A \leq A \vee B$ and $B \leq A \vee B$.
- Whenever $A \leq X$ and $B \leq X$ for an element $X \in L$, then also $A \vee B \leq X$.

Dually, for $A, B \in L$ the element $A \wedge B \in L$ satisfies:

- $A \wedge B \leq A$ and $A \wedge B \leq B$.
- Whenever $Y \leq A$ and $Y \leq B$ for an element $Y \in L$, then also $Y \leq A \wedge B$.

For this reason, these operations are sometimes called the *least upper bound* and the *greatest lower bound* respectively.

A *ranked lattice* is a pair (L, r), where L is a lattice and the *rank function* $r : L \to \mathsf{N}$ satisfies the property that $r(Y) = r(X) + 1$ whenever Y *covers* X in L, that is, $X \leq Y$ and if $X \leq Z \leq Y$ then $Z = X$ or $Z = Y$. One can consider the natural generalisation of the inclusion–exclusion formula (1.2) in the setting of a general ranked lattice. One can ask: For which ranked lattices does the following generalised inclusion-exclusion formula hold?

$$r(\bigvee_{i\in[n]} A_i) = \sum_{\emptyset\neq I\subseteq[n]} (-1)^{|I|-1} r(\bigwedge_{i\in I} A_i). \tag{2.3}$$

The case $n = 2$ of this formula,

$$r(A) + r(B) = r(A \vee B) + r(A \wedge B), \tag{2.4}$$

is obviously a necessay condition for the general formula to hold. When (2.4) holds, the rank function r is said to be *modular*.

We give now some natural examples of ranked lattices that are important in combinatorics. Let's see if the generalised inclusion–exclusion formula holds in these cases.

Example 10 (Sets) The first example is our starting point: the lattice of subsets of a universal set U. The lattice is $(2^U, \cup, \cap, \emptyset, U)$, and the rank function is just the cardinality. The fact that the inclusion–exclusion formula holds is one of the first theorems in any textbook on combinatorics.

Example 11 (Divisors) Let N be a positive integer and let $\mathcal{D}_N$ denote the set of all positive divisors of N, ordered by divisibility i.e. for divisors a, b of N, $a \leq b$ iff a divides b. Let $N := p_1^{\alpha_1} \cdot \cdots \cdot p_r^{\alpha_r}$ be the (unique) prime factorisation of N. The lattice operations can be described by regarding an element a as a r vector $(a_1, \ldots, a_r)$ where a_i is aan integer with $0 \leq a_i \leq \alpha_i$ for each $i \in [r]$. The elements $a \vee b$ and $a \wedge b$ are given by the following equations on the components $i \in [r]$:

$$(a \vee b)_i := \max(a_i, b_i) \quad (a \wedge b)_i := \min(a_i, b_i).$$

The bottom element is 1 and the top element is N.

The rank of an element $a := (a_1, \ldots, a_r)$ is $\sum_i a_i$. The inclusion–exclusion formula holds in this ranked lattice.

Example 12 (Multisets) This example generalises the previous one. Let $\mathbf{e} := (e_1, e_2, \ldots)$ be an (infinite) vector whose components are non–negative integers or ∞. Let $\mathcal{M}_\mathbf{e}$ denote the collection of "all finite multisets of integers with multiplicities restricted to $\mathbf{e}$". By this we mean the family of all finite unordered collections of positive integers, where repetitions are allowed but each i can appear at most e_i times. Each multiset in $\mathcal{M}_\mathbf{e}$ can be represented by a vector $\sigma := (\sigma_1, \sigma_2, \ldots)$ of non–negative integers, where $0 \leq \sigma_i \leq e_i$ for each $i \geq 1$, and also $\sum_i \sigma_i < \infty$. This collection forms a partial order under component–wise ordering. In fact it forms a lattice, where the lattice operations are given by the max and min operations respectively, as in the previous example. The bottom element is $(0, 0, , \ldots)$ and the top element is $\mathbf{e}$.

The rank of an element σ in this lattice is $\sum_i \sigma_i$. The inclusion–exclusion principle holds in this ranked lattice.

Example 13 (Partitions) Let m be a positive integer and let $\mathcal{P}_m$ denote the set of all *partitions* of the set $[m]$ ordered by (inverse) *refinement*. Thus $\Pi_1 \leq \Pi_2$ if the partition Π_1 refines the partition Π_2, that is, each block of the partition Π_2 is the union of some blocks of partition Π_1. The bottom element is the partition into singletons and the top element is the partition in which all elements are in a single block. The join of partitions Π_1 and Π_2, is the partition which groups two elements together in a block if they are in the same block in one of the two partitions and the meet is the partition that groups two elements together in a block if they are in the same block in both partitions.

The rank of a partition Π is $n - k$ where k is the number of blocks in Π. The inclusion–exclusion formula fails in this ranked lattice. To see this, take the lattice $\mathcal{P}_3$ and check the formula for $n = 3$ with the partitions $\{23, 1\}, \{31, 2\}$ and $\{12, 3\}$.

Example 14 (Vector Spaces) Let V be a finite–dimensional vector space over a field k and let $L(V,k)$ denote the family of all k–subspaces of V ordered by the subspace relation. The meet operation is just set intersection. The join of the subspaces U and W is given by the *sum* $U + W := \{u + w \mid u \in U, w \in W\}$; this is the least subspace that contains both U and W. (Note that the set theoretic union of two subspaces is not necessarily a subspace!) The bottom element is the trivial space $\{\mathbf{0}\}$ and the top element is V.

The rank of a subspace W is its dimension, dim W. The inclusion–exclusion formula does *not* hold in general in this ranked lattice. To see this, consider the plane, R^2 regarded as a vector space over the reals. Let X be the one dimensional space spanned by $(1,0)$, namely the x axis, let Y be the one dimensional space spanned by $(0,1)$, namely the y axis, and let Z be the one dimensional space spanned by $(1/2,1/2)$. Then, $X+Y+Z$ is the whole plane, and has dimension 2. However, each pairwise intersection is trivial, so the right–hand side of the inclusion–exclusion formula gives 3 and not 2.

What goes wrong in the last two examples? Can the fault be with the rank function? One can check that the rank function in each case is indeed modular, so the fault lies elsewhere (in our star?!). Let's discover where it is, by trying to "prove" (2.3).

We proceed by induction on n. For $n = 1$ this is trivial and for $n = 2$, we have the modular law. For $n \geq 3$, we compute,

$$
\begin{aligned}
r(\bigvee_{i\in[n]} A_i) &= r(\bigvee_{i\in[n-1]} A_i \vee A_n) \\
&= r(\bigvee_{i\in[n-1]} A_i) + r(A_n) - r((\bigvee_{i\in[n-1]} A_i) \wedge A_n), \quad \text{using the modular law} \\
&= r(\bigvee_{i\in[n-1]} A_i) + r(A_n) - r(\bigvee_{i\in[n-1]} (A_i \wedge A_n)), \quad (*) \\
&= \sum_{\emptyset\neq I\subseteq[n-1]} (-1)^{|I|-1} r(\bigwedge_{i\in I} A_i) + r(A_n) + \sum_{\emptyset\neq J\subseteq[n-1]} (-1)^{|J|-1} r(\bigwedge_{j\in J}(A_j \wedge A_n)), \\
&\qquad \text{by induction} \\
&= \sum_{\emptyset\neq I\subseteq[n],n\notin I} (-1)^{|I|-1} r(\bigwedge_{i\in I} A_i) + \sum_{\emptyset\neq I\subseteq[n],n\in I} (-1)^{|I|-1} r(\bigwedge_{i\in I} A_i) \\
&= \sum_{\emptyset\neq I\subseteq[n]} (-1)^{|I|-1} r(\bigwedge_{i\in I} A_i).
\end{aligned}
$$

And the generalised inclusion–formula holds! But wait a minute, something must be wrong, because we *know*, from our examples, that it doesn't hold in general! Checking our derivation again, we see that in the line labelled $(*)$, we cheated. This step is valid only if the lattice is *distributive* i.e. if for all $A, B, C \in L$,

$$(A \vee B) \wedge C = (A \wedge C) \vee (B \wedge C).$$

However, if the lattice is distributive, then we have a perfectly valid proof of the the general inclusion–exclusion formula. Incidentally, note that this property of distributivity indeed distinguishes the two classes of examples before. Also, as a by–product, we note that the lattice of subsets is distributive, and hence finally we have a proof of our original formula (1.2) (which we had kept on deferring).

What about a converse? Suppose we know that the generalised inclusion exclusion formula (2.3) holds in a ranked lattice (L, r). Applying the formula with $n = 3$ and $A_3 := \bot$, we swiftly recover the modular law. Can we also get distributivity?

Let's compute $r(A \vee B \vee C)$ for arbitrary $A, B, C \in L$. Applying the modular law twice, we get for any $A, B, C \in L$:

$$\begin{aligned} r(A \vee B \vee C) &= r(A \vee B) + r(C) - r((A \vee B) \wedge C) \\ &= r(A) + R(B) + r(C) - r(A \wedge B) - r((A \vee B) \wedge C) \end{aligned} \tag{2.5}$$

Also, applying the inclusion–exclusion formula for $n = 3$ directly and then the modular law gives

$$\begin{aligned} r(A \vee B \vee C) &= r(A) + r(B) + r(C) - r(B \wedge C) - r(C \wedge A) - r(A \wedge B) + r(A \wedge B \wedge C) \\ &= r(A) + r(B) + r(C) - r((B \wedge C) \vee (C \wedge A)) - r(A \wedge B). \end{aligned} \tag{2.6}$$

Comparing equations (2.5) and (2.6), we see that

$$r((A \vee B) \wedge C) = r((B \wedge C) \vee (C \wedge A)). \tag{2.7}$$

Since $(B \wedge C) \vee (C \wedge A) \leq (A \vee B)) \wedge C$, (why?) equation (2.7) implies that in fact $(B \wedge C) \vee (C \wedge A) = (A \vee B)) \wedge C$, that is, the lattice is distributive!

All in all, we've shown the following cycle of implications:

Theorem 15 *Let (L, r) be a ranked lattice. The following are equivalent:*

1. *L is a distributive lattice and r is modular.*
2. *The general inclusion–exclusion formula (2.3) for every positive integer n holds for (L, r).*
3. *The inclusion–exclusion formula (2.3) for $n = 3$ holds for (L, r).*

The implication (3) $\rightarrow$ (2) is a bit surprising and I haven't found it mentioned in the literature.

Chapter 3

Möbius Inversion

We will now develop a technique that will enable us to directly deduce the value of the Euler totient function ϕ, by *inverting* (2.1),

$$n = \sum_{d|n} \phi(d),$$

In fact, the technique will be general enough to invert the following types of relations between functions $f, g, h : \mathsf{L} \to \mathsf{R}$ over an *arbitrary* poset P.

$$g(x) = \sum_{y \leq_P x} f(y), \tag{3.1}$$

$$h(x) = \sum_{y \geq_P x} f(y) \tag{3.2}$$

In fact, once we set up the appropriate machinery, this will be a piece of cake!

Let P be an arbitrary poset. The *Zeta-function* of P, $\zeta_P : P \times P \to \{0, 1\}$ is then defined as

$$\zeta_P(x, y) := \begin{cases} 1, & \text{if } x \leq y; \\ 0, & \text{otherwise.} \end{cases} \tag{3.3}$$

or, in our handy notation,

$$\zeta_P(x, y) := [x \leq y]. \tag{3.4}$$

When the poset is clear from the context, we'll omit the subscript.

Example 16 *Consider the lattice $\mathcal{D}_{12}$ of divisors of 12, shown in the following figure:*

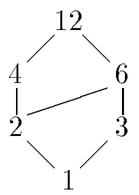

Figure 3.1: The lattice $\mathcal{D}_{12}$.

We can represent $\zeta_{\mathcal{D}_{12}}$ as a matrix indexed by the elements of $\mathcal{D}_{12}$; note that e.g. $\zeta_{\mathcal{D}_{12}}(2,3) = 0$, since $2 \not\leq 3$ in the lattice $\mathcal{D}_{12}$!

$$\zeta_{\mathcal{D}_{12}} := \begin{array}{c} \\ 2 \\ 1 \\ 3 \\ 12 \\ 6 \\ 4 \end{array} \begin{array}{c} \begin{array}{cccccc} 2 & 1 & 3 & 12 & 6 & 4 \end{array} \\ \boxed{\begin{array}{cccccc} 1 & 0 & 0 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 \end{array}} \end{array} \quad \text{or} \quad \begin{array}{c} \\ 1 \\ 2 \\ 3 \\ 4 \\ 6 \\ 12 \end{array} \begin{array}{c} \begin{array}{cccccc} 1 & 2 & 3 & 4 & 6 & 12 \end{array} \\ \boxed{\begin{array}{cccccc} 1 & 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{array}} \end{array}$$

Notice that the "naturally indexed" matrix (on the right hand side) is upper triangular.

The fact that the right hand matrix is upper triangular is not an accident. The elements of *any* finite poset P can be listed in such a way that the Zeta matrix of P is upper–triangular. This is particularly easy for us as computer scientists to see – simply draw the diagram of the poset as above [1] and number it bottom–up in a breadth–first manner! Alternatively, we can achieve this inductively as follows: delete a maximal element, number the remaining poset, and put the deleted element back, numbering it last.

Let $|P| = n$, and let $\{x_1, x_2, \ldots, x_n\}$ be a numbering of the elements of P, such that whenever $x_i \leq_P x_j$, then $i \leq j$. With this, it is easy to show by induction on the size of the lattices that ζ_P is upper triangular. In the example above, the "natural" numbering $x_1 := 1, x_2 := 2, x_3 := 3, x_a := 4, x_5 := 6, x_7 := 12$ has this property.

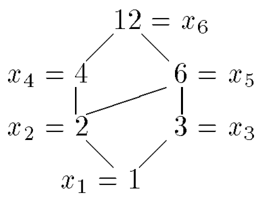

Figure 3.2: L_{12} with numbering.

The set T, of all such upper triangular matrices indexed by elements of the poset has nice properties. First, it is a R–*algebra.* What this means is that

- If $A \in T$ and $c \in \mathsf{R}$, then $c \cdot A \in T$.
- If $A, B \in T$, then also $A + B \in T$.
- If $A, B \in T$, then also $AB \in T$.
- If $A \in T$ is *invertible* over the reals, then $A^{-1} \in T$.

This algebra is called the *incidence algebra* of the poset.

Another extremely nice property of upper triangular matrices is that one can compute their determinants easily: simply multiply the elements on the diagonal! For example, $\det\zeta = 1$.

Let's recall how to multiply two matrices are multiplied, just in case we're a bit rusty in that department. The product of two matrices $C = AB$, is a matrix defined when the

[1] This is called the *Hasse Diagram* of the poset.

matrices A and B are compatible i.e. the number of columns of A is equal to the number of rows of B. In our case, we always deal with *square* matices of size n, where n is the number of elements in P. The matrix C has elements given by the formula:

$$C_{ij} = \sum_{1 \leq k \leq n} A_{ik} B_{kj},$$

which we in the following will write in functional notation and dropping the subscript range as per our conventions as:

$$C(i,j) = \sum_k A(i,k)B(k,j)$$

Just to test out our matrix multiplication skills, let's compute ζ^2:

$$\begin{aligned}
\zeta^2(x,y) &= \sum_z \zeta(x,z)\zeta(z,y) \\
&= \sum_z [x \leq z][z \leq y] \\
&= \sum_{x \leq z \leq y} 1 \\
&= \#\{z \mid x \leq z \leq y\}
\end{aligned}$$

A configuration $x < z < y$ in the poset is called a *chain* of *length* 1. Let's continue and

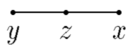

Figure 3.3: A chain of length 2

compute ζ^3:

$$\begin{aligned}
\zeta^3(x,y) &= \sum_z \zeta(x,z)\zeta^2(z,y) \\
&= \sum_z [x \leq z] \sum_{z \leq z' \leq y} 1 \\
&= \sum_{x \leq z} 1 \sum_{z \leq z' \leq y} 1 \\
&= \sum_{x \leq z \leq z' \leq y} 1
\end{aligned}$$

A configuration $x < z < z' < y$ is a chain of length 2: In general, by induction we can

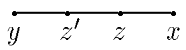

Figure 3.4: A chain of length 3

compute ζ^k for any $k \geq 0$ and verify that

$$\zeta^k = \sum_{x \leq z_1 \leq z_2 \leq \ldots \leq z_{k-1} \leq y} 1.$$

A configuration $x := z_0 < z_1 < z_2 < \cdots < z_{k-1} < z_k =: y$ is called a chain of lenghth k. Our functions are computing the number of such chains ... well, almost. There's the irritant that we have $\leq$ rather than $<$ in our equations, so we're not quite counting the right things. But we'll set it right as follows.

Let's look at the function given by $(\zeta - 1)$, (1 denoting the $(n \times n)$ identity matrix I_n with ones on the diagonal entries and zeroes elsewhere). Clearly, the resulting upper triangular representation has ones on entries (i, j) where $x_i < x_j$, i.e.

$$(\zeta - 1)(x, y) = \left\{ \begin{array}{ll} 0 & x \geq y \\ 1 & x < y \end{array} \right\},$$

or,

$$(\zeta - 1)(x, y) = [x < y].$$

Exercise 17 *Prove that*

$$\begin{aligned} (\zeta - 1)^k(x, y) &= \sum_{x := z_0 < z_1 < \cdots < z_{k-1} < z_k =: y} 1 \\ &= \#\{ \textit{ chains of length } k \textit{ between } x \textit{ and } y \ \}. \end{aligned}$$

As a last calculation, we want to determine the function $(2 - \zeta)^{-1}(x, y)$. Of course, we must ensure that $(2 - \zeta)$ is invertible. But that's easy: take a look at the diagonal, it has only ones, hence the determinant is 1 and the matrix is invertible!

To find $(2-\zeta)^{-1}$, we shall use the following black magic; it will give you a taste of "Things to Come" [2]:

$$\frac{1}{1 - z} = 1 + z + z^2 + z^3 + \cdots = \sum_{k \geq 0} z^k.$$

Hence:

$$(2 - \zeta)^{-1} = \frac{1}{1 - (\zeta - 1)} = 1 + (\zeta - 1) + (\zeta - 1)^2 + \cdots = \sum_{k \geq 0} (\zeta - 1)^k.$$

and so:

$$\begin{aligned} (2 - \zeta)^{-1}(x, y) &= \left(\sum_{k \geq 0} (\zeta - 1)^k \right) (x, y) \\ &= \sum_{k \geq 0} (\zeta - 1)^k (x, y) \\ &= \sum_{k \geq 0} \#\{ \text{ chains of length } k \text{ between } x \text{ and } y \ \} \\ &= \#\{ \text{ chains between } x \text{ and } y \ \} \end{aligned}$$

Let's stop fooling around and return to the main business at hand. Go back to (3.1) and (3.2). Let us regard a function $f : \mathsf{P} \to \mathsf{R}$ as a vector indexed by the elements of P:

$$f := (f(x_1), f(x_2), \ldots, f(x_n))^T.$$

[2] Title of a lovely jig by Dizzy Gillespie.

Now let us applying the ζ-function to f:

$$\begin{aligned}(\zeta f)(x) &= \sum_y \zeta(x,y)f(y)\\ &= \sum_y [x \le y] f(y)\\ &= \sum_{x\le y} f(y)\\ &= h(x).\end{aligned}$$

and,

$$\begin{aligned}(\zeta^T f)(x) &= \sum_y \zeta^T(x,y)f(y)\\ &= \sum_y \zeta(y,x)f(y)\\ &= \sum_{y\le x} f(y)\\ &= g(x).\end{aligned}$$

In matrix notation, (3.1) and (3.2) take the compact form:

$$g = \zeta^T \cdot f \quad \text{and} \quad h = \zeta f. \tag{3.5}$$

This is the promised piece of cake: the solution to this set of equations is trivial:

$$f = (\zeta^{-1})^T \cdot g \quad \text{and} \quad f = \zeta^{-1} \cdot h \tag{3.6}$$

Let us introduce a special name for ζ^{-1}:

$$\mu := \zeta^{-1}. \tag{3.7}$$

This is called the *Möbius function* of the poset. Using this, we can rewrite (3.6) as follows:

$$f = \mu^T \cdot g \quad \text{and} \quad f = \mu \cdot h \tag{3.8}$$

When expanded out, the solutions to (3.1) and (3.2) are:

$$f(x) = \sum_{y\le x} \mu(y,x)g(y), \tag{3.9}$$

and

$$f(x) = \sum_{x\le y} \mu(x,y)h(y), \tag{3.10}$$

But it is easiest to remember the matrix forms. Table 3.1 summarizes the results.

Thus, the Möbius-function is the key to the inversion problem: once we know the Möbius function of a poset, we're in business to perform the inversion.

$$\boxed{h = \zeta f}, \boxed{f = \mu h},$$
$$\boxed{g = \zeta^T f}, \boxed{f = \mu^T g}.$$

Table 3.1: Möbius Inversion formulae.

By definition, $\zeta\mu = 1 = \mu\zeta$. Thus,

$$\begin{aligned}[x = y] &= (\zeta\mu)(x, y) \\ &= \sum_z \zeta(x, z)\mu(z, y) \\ &= \sum_{x \leq z \leq y} \mu(z, y)\end{aligned}$$

Taking $y = x$, we see that for all x, $\mu(x, x) = 1$; that is, the diagonal entries of the matrix μ are always 1. Next writing the last equation as

$$\mu(x, y) + \sum_{x \leq z < y} \mu(z, y) = 0$$

gives us a method to determine the values of the Möbius-function in a "top–down" manner, namely

$$\mu(x, y) = -\sum_{x \leq z < y} \mu(z, y). \tag{3.11}$$

Alternatively,

$$\begin{aligned}[x = y] &= (\mu\zeta)(x, y) \\ &= \sum_z \mu(z, y)\zeta(x, z) \\ &= \sum_{x \leq z \leq y} \mu(x, z)\end{aligned}$$

can be used for a "bottom-up" calculation of the values of μ:

$$\mu(x, y) = -\sum_{x \leq z < y} \mu(x, z). \tag{3.12}$$

Let us try it out on a simple example. Consider the lattice $\mathsf{L} = (\{1, 2, \ldots, n\}, \leq)$ where $\leq$ is the natural ordering on the integers. Let us calculate the Möbius function. First, since mu is upper triangular, we know that $\mu(x, y) = 0$ for $y < x$. Also, we know that $\mu(x, x) = 1$ always. So we are only left with computing $\mu(x, y)$ for $x < y$. Let's try the simplest case: $y = x + 1$. We have, by (3.11) or (3.12),

$$\mu(x, x + 1) = -\mu(x, x) = -1.$$

Next, we try $\mu(x, x + 2)$ and again by (3.11) or (3.12),

$$\mu(x, x + 2) = -(\mu(x, x) + \mu(x, x + 1)) = 0.$$

We can continue this, and for any $k \geq 0$, we have by (3.11) or (3.12),

$$\mu(x, x+k) = -(\mu(x,x) + \mu(x, x+1) + \cdots + \mu(x, x+k-1).$$

This means of course, that $\mu(x, x+k) = 0$ for all $k > 1$. We have fully determined the Möbius function of the chain! To summarise in our shorthand notation,

$$\mu(x, y) = [x = y] - [y = x + 1]. \tag{3.13}$$

Now that we have the Möbius function, we can't resist putting it to use: so if

$$g(x) := \sum_{y \leq x} f(y),$$

what is the inverse relation? Applying Table 3 or (3.9), and substituting the Möbius function we just determined, we get:

$$\begin{aligned} f(x) &= \sum_{y \leq x} \mu(y, x) g(y) \\ &= \sum_{y \leq x} ([x = y] - [x = y + 1]) g(y) \\ &= g(x) - g(x-1). \end{aligned}$$

Not very exciting! But at least it's reassuring.

Exercise 18 *Let* P *be a (finite) poset with a unique bottom element* $\perp$ *and a unique top element,* $\top$*. For* $i \geq 0$*, denote by* c_i *the number of chains* $\perp = x_0 < x_1 < \cdots < x_i = \top$ *of length* i *between* $\perp$ *and* $\top$*. (Thus* $c_0 = 0$ *and* $c_1 = 1$*.) Show that*

$$\mu_P(\perp, \top) = \sum_i (-1)^i c_i.$$

Exercise 19 *Let* P *be a forest i.e. for each* $a \in P$*, the elements* $\{x \in P \mid x \leq a\}$ *form a chain. (Our forests grow upwards, as they should!). Compute the Möbius function of* P.

It would be nice to compute the Möbius functions of more complicated posets. Here is a standard construction to build more complex posets from simple ones:

Definition 20 *Let* $\mathsf{P}_1 := (P_1, \leq_1)$ *and* $\mathsf{P}_2 := (P_2, \leq_2)$ *be two partially ordered sets. The* product *of* P_1 *and* P_2 *is*

$$\mathsf{P}_1 \times \mathsf{P}_2 := (P_1 \times P_2, \leq),$$

where

$$(x_1, x_2) \leq (y_1, y_2) \quad \Longleftrightarrow \quad x_1 \leq_1 x_2 \text{ and } y_1 \leq_2 y_2.$$

Suppose we know the Möbius functions of P_1 and P_2. How can we use that to compute the Möbius function of their product? To answer this, let's look at the Zeta function:

$$\begin{aligned} \zeta((x_1, x_2), (y_1, y_2)) &= [(x_1, x_2) \leq (y_1, y_2] \\ &= [x_1 \leq_1 y_1][x_2 \leq_2 y_2] \\ &= \zeta_1(x_1, y_1) \cdot \zeta_2(x_2, y_2). \end{aligned}$$

Motivated by this, we suspect that the same formula should also work with μ:

$$\mu_{P_1 \times P_2}((x_1, x_2), (y_1, y_2)) = \mu_{P_1}(x_1, y_1) \cdot \mu_{P_2}(x_2, y_2). \tag{3.14}$$

To confirm this, we compute:

$$\begin{aligned}
\zeta\mu((x_1, x_2), (y_1, y_2)) &= \sum_{(u,v)} \zeta((x_1, x_2), (u, v)) \cdot \mu((u, v) \cdot (y_1, y_2)) \\
&= \sum_{(u,v)} \zeta_1(x_1, u)\mu_1(u, y_1)\zeta_2(x_2, v)\mu_2(v, y_2) \\
&= \sum_{u} \zeta_1(x_1, u)\mu_1(u, y_1) \sum_{v} \zeta_2(x_2, v)\mu_2(v, y_2) \\
&= [x_1 = y_1][x_2 = y_2] \\
&= [(x_1, x_2) = (y_1, y_2)].
\end{aligned}$$

Which shows that indeed we have the right formula!

The operation on matrices in (3.14) is a funny kind of matrix product. It is called the *Kronecker* or *tensor* product of matrices. If A is a $m \times n$ matrix and B is a $p \times q$ matrix, then their tensor product, $A \otimes B$ is the $mp \times nq$ matrix given by:

$$A \otimes B := \left[\begin{array}{c|c|c|c}
a_{11}B & a_{12}B & \cdots & a_{1,n}B \\ \hline
a_{21}B & a_{22}B & \cdots & a_{2,n}B \\ \hline
\vdots & \vdots & \vdots & \vdots \\ \hline
a_{m1}B & a_{m,2}B & \cdots & a_{m,n}B
\end{array}\right].$$

Hence, in matrix notation, we can write:

$$\mu_{P_1 \times P_2} = \mu_{P_1} \otimes \mu_{P_2}. \tag{3.15}$$

Exercise 21 *Verify that $(A \otimes B)^T = A^T \otimes B^T$ and that if A and B are non-singular, then $(A \otimes B)^{-1} = A^{-1} \otimes B^{-1}$. Hence deduce the formula (3.14) directly.*

To start with a simple example, consider the *lattice of truth values*

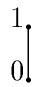

This one's Möbius-function is given by

$$\mu_2 \quad := \quad \begin{array}{c|cc|}
 & 0 & 1 \\ \hline
0 & 1 & -1 \\
1 & 0 & 1 \\ \hline
\end{array}$$

Now, let us form the product of two of these lattices; this yields the Boolean algebra of 4 elements:

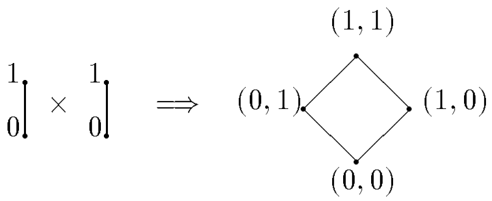

To compute the Möbius function, we compute the tensor product:

$$\begin{array}{c|cc|} & 0 & 1 \\ \hline 0 & 1 & -1 \\ 1 & 0 & 1 \\ \hline \end{array} \otimes \begin{array}{c|cc|} & 0 & 1 \\ \hline 0 & 1 & -1 \\ 1 & 0 & 1 \\ \hline \end{array} = \begin{bmatrix} 1 & -1 & -1 & 1 \\ 0 & 1 & 0 & -1 \\ 0 & 0 & 1 & -1 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

We can iterate this many times – and to a purpose!

Let n be a positive integer and consider the lattice $((2^{[n]}, \cup, \cap, \emptyset, [n])$. We can identify a subset $I \subseteq [n]$ with its *characteristic* vector $(i_1,i_2,\ldots,i_n)$ in 2^n such that

$$i_k :== \begin{cases} 1, & \text{if } k \in I; \\ 0, & \text{otherwise.} \end{cases}$$

Let $I \subseteq J \subseteq [n]$. If we identify J with its characteristic vector $j_1,j_2,\ldots,j_n$, then notice first that $i_k \leq j_k$ for all $k \in [n]$. Let us compute: using the tensor product formula (3.15) in the second line:

$$\begin{aligned} \mu(I,J) &= \mu(\,(i_1,i_2,\ldots,i_n)\,,\,(j_1,j_2,\ldots,j_n)\,) \\ &= \mu(i_1,j_1)\mu(i_2,j_2)\cdots\mu(i_n,j_n) \\ &= (-1)^{j_1-i_1}(-1)^{j_2-i_2}\cdots(-1)^{j_n-i_n} \\ &= (-1)^{\sum_k (j_k-i_k)} \\ &= (-1)^{|J|-|I|}. \end{aligned} \tag{3.16}$$

In the second line, we used the tensor product formula (3.15), and in the third line, we used the values of μ_2.

Now, take a family $A_1,A_2,\ldots,A_n$ of subsets of a universe U, and for an index set $I \subseteq [n]$, set

$$h(I) := \left|\bigcap_{i\in I} A_i\right|, \qquad f(I) := \left|\bigcap_{i\in I} A_i \cap \bigcap_{i\notin I} \bar{A}_i\right|.$$

Observe that $h(I) = \sum_{J\supseteq I} f(J)$. Inverting this using (3.10) we get, using (3.16) in the third equality,:

$$\begin{aligned} f(I) &= \sum_J \mu(I,J)h(J) \\ &= \sum_{I\subseteq J} \mu(I,J)h(J) \\ &= \sum_{I\subseteq J} (-1)^{|J|-|I|}h(J), \end{aligned}$$

That is,

$$\left|\bigcap_{i\in I} A_i \cap \bigcap_{i\notin I} \bar{A}_i\right| = \sum_{I\subseteq J} (-1)^{|J|-|I|} \left|\bigcap_{j\in J} A_j\right|.$$

Setting $I = \emptyset$, we have:

$$|\bar{A}_1 \cap \bar{A}_2 \cap \ldots \cap \bar{A}_n| = \sum_J (-1)^{|J|} \left| \bigcap_{j \in J} A_j \right|,$$

and using de Morgan's law, this is equivalent to:

$$|U| - |A_1 \cup A_2 \cup \ldots \cup A_n| = \left| \bigcap_{j \in \emptyset} A_j \right| + \sum_{\emptyset \neq J \subseteq [n]} (-1)^{|J|} \left| \bigcap_{j \in J} A_j \right|.$$

This may seem a bit mysterious; but notice that $\bigcap_{j \in \emptyset} A_j = U$, hence, we finally have:

$$|A_1 \cup A_2 \cup \ldots \cup A_n| = \sum_{\emptyset \neq J \subseteq [n]} (-1)^{|J|-1} \left| \bigcap_{j \in J} A_j \right|.$$

We have recovered the inclusion–exclusion formula from the first lecture! Thus we have seen that Möbius inversion yields the principle of inclusion–exclusion as a special case.

Exercise 22 *In this exercise, we develop the* product form *of the Möbius inversion formula. Let*

$$g(x) = \prod_{y \leq x} f(y).$$

Show that then

$$f(x) = \prod_{y \leq x} g(y)^{\mu(y,x)}.$$

Exercise 23 *Let m be a positive integer and let*

$$\omega := e^{2\pi i/m} = \cos(2\pi/m) + i\sin(2\pi/m).$$

We say that ω is an mth root of unity *since $\omega^m = e^{2\pi i} = 1$, but $\omega^k \neq 1$ for $k < m$. In fact, each of the m complex numbers $\omega^0, \omega^1, \ldots, \omega^{m-1}$ is also an mth root of unity. Hence $z - \omega^k$ is a factor of the polynomial $z^m - 1$ for each $0 \leq k < m$. Since these factors are distinct, the complete factorisation of $z^m - 1$ over the complexes is:*

$$z^m - 1 = \prod_{0 \leq k < m} (z - \omega^k).$$

1. *Let $\Psi_m(z) := \prod_{0 \leq k < m, (k,m)=1} (z - \omega^k)$. This is called the* cyclotomic polynomial of order *m. Prove that*

$$z^m - 1 = \prod_{d|m} \Psi_d(z).$$

2. *Invert this using the previous exercise to show that*

$$\Psi_m(z) = \prod_{d|m} (z^d - 1)^{\mu(m/d)}.$$

Chapter 4

Möbius Inversion II

In this lecture, we'll continue where we left off and compute the Möbius functions for a few more lattices of interest.

For a start, we are now ready to use the technique presented in the last chapter to invert the equation

$$n = \sum_{d|n} \phi(d),$$

to get the value of the Euler totient function ϕ. Remembering the ordering on the lattice of divisors $\mathcal{D}$ we see that this is an example of the form $g(n) = \sum_{d \leq n} f(d)$ $(g = \zeta^T f)$ where g is the identity function id and f is the Euler totient function. So we have $\phi = \mu^T id$.

The job then is to compute the μ-function for a divisor lattice $\mathcal{D}_n$. To tackle this, we observe that a divisor lattice $\mathcal{D}_n$ is isomorphic to the product of a number of linear ordered lattices. To wit, if $N = p_1^{\alpha_1} p_2^{\alpha_2} \ldots p_r^{\alpha_r}$ every element in the Divisor lattice $\mathcal{D}_n$ can be represented as an r-tuple: $(\beta_1, \beta_2, \ldots, \beta_r)$, where $0 \leq \beta_i \leq \alpha_i$, for $1 \leq i \leq r$ i.e. each entry is the exponent of the corresponding prime factor. This gives us exactly the product lattice

$$\mathsf{C}_1 \times \mathsf{C}_2 \times \ldots \mathsf{C}_r, \quad \text{where } \mathsf{C}_i \text{ is the linear ordered } \textit{chain} \text{ lattice } \{0, 1, \ldots, \alpha_i\}.$$

Example 24 *Consider the lattice L_{12} pictured in the last section (figure 3.1). We have $12 = 2^2 * 3^1$ so the isomorphic lattice is $\{0,1,2\} \times \{0,1\}$ and the relabelling goes as:*

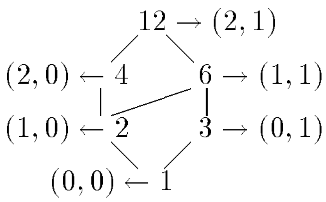

Figure 4.1: L_{12} with relabeling.

We are now in position to calculate the μ-function for the divisor lattice, since the μ-function is the same for isomorphic lattices. From the previous chapter we know how to compute the μ-function on the linear ordered chains (3.13) and we are also armed with the product formula (3.14). Recall that for a chain, $\mu(x,y) = 0$ unless $0 \leq y - x \leq 1$ and then the value is $(-1)^{y-x}$. Suppose $a \leq_d b$, where $a := (a_1, a_2, \ldots, a_r)$, and $b :== (b_1, b_2, \ldots, b_r)$

are the product representations (note that therefor, $a_i \le b_i$ for each i). Then:

$$\begin{aligned}
\mu(a,b) &= \mu((a_1,a_2,\ldots,a_r),(b_1,b_2,\ldots,b_r)) \\
&= \mu(a_1,b_1)\mu(a_2,b_2)\ldots\mu(a_r,b_r) \\
&= \begin{cases} (-1)^{\sum_i (b_i-a_i)}, & \text{if} \quad b_i - a_i \le 1; \\ 0 & \text{othewise.} \end{cases}
\end{aligned} \tag{4.1}$$

In the above example we could compute $\mu(2,4) = \mu((1,0),(2,0)) = (-1)^1 = -1$, and $\mu(1,4) = \mu((0,0),(2,0)) = 0$

Observe that $\mu(2,12) = \mu(1,6)$ and $\mu(4,12) = \mu(2,6)$. In general, we will use the following fact repeatedly.

Fact 25 *The value of $\mu(A,B)$ in a lattice L where $A \le_L B$ only depends on the structure of the sublattice of L between A and B. We denote this sublattice $\mathsf{L}_{|(A,B)}$.*

Having μ explicitly we now get from $\phi = \mu^T id$:

$$\begin{aligned}
\phi(n) &= \sum_{d|n} \mu^T(n,d)\cdot d \\
&= \sum_{d|n} \mu(d,n)\cdot d \\
&= \sum_{d|n} \mu(1,\frac{n}{d})\cdot d \\
&= \sum_{d'|n} \mu(1,d')\cdot(\frac{n}{d'}) \\
&= \sum_{d|n} \mu(d)\cdot(\frac{n}{d}) \quad \text{using } \mu(d) \text{ as short for } \mu(1,d)
\end{aligned}$$

The equality $\mu(d,n) = \mu(1,\frac{n}{d})$ follows if you consider (4.1),(dividing with d does not change $(b_i - a_i)$ for any i).

The second last equation is just doing the summation in a different order.

At last, using (4.1) we get

$$\mu(d) = \begin{cases} (-1)^r & \text{if } d = p_1p_2\ldots p_r \text{ for distinct primes;} \\ 1 & \text{if } d = 1 \text{ ;} \\ 0 & \text{otherwise.} \end{cases}$$

To give an example, we can compute:

$$\begin{aligned}
\phi(12) &= \mu(1)\cdot 12 + \mu(2)\cdot 6 + \mu(3)\cdot 4 + \mu(6)\cdot 2 \\
&= 12 - 6 - 4 + 2 \\
&= 4
\end{aligned}$$

corresponding to the 4 numbers $1,5,7,11 \in [n]$ which are relatively prime to 12.

If want to compute μ for a non-distributive lattice, we can't hope to find an isomorphic lattice which is a product of linear ordered sets. This follows from the fact, that linear ordered sets are distributive and forming the product preserves distributivity.

Exercise 26 *Verify this assertion!*

Instead we can use Weisner's theorem:

Theorem 27 (Wiesner) *Let* $\mathsf{L} = (L, \vee, \wedge, \bot, \top)$ *be a finite lattice and let* $\alpha \in \mathsf{L}\quad, \alpha > \bot$*. Then*

$$\sum_{x: x\vee\alpha=\top} \mu(\bot, x) = 0$$

Proof:

$$\begin{aligned}
\sum_{x: x\vee\alpha=\top} \mu(\bot, x) &= \sum_{x} \mu(\bot, x)[x \vee \alpha = \top] \\
&= \sum_{x} \mu(\bot, x) \sum_{y: x\vee\alpha \le y \le \top} \mu(y, \top) \qquad ([a = b] = (\zeta\mu)(a, b)) \\
&= \sum_{x} \mu(\bot, x) \sum_{y \ge x, y \ge \alpha} \mu(y, \top) \\
&= \sum_{y: y\ge\alpha} \sum_{x: x\le y} \mu(\bot, x) \cdot \mu(y, \top) \\
&= \sum_{y: y\ge\alpha} \mu(y, \top) \sum_{x: x\le y} \mu(\bot, x) \\
&= \sum_{y: y\ge\alpha} \mu(y, \top) \cdot [\bot = y] \qquad ([a = b] = (\mu\zeta)(a, b)) \\
&= 0 \qquad (\text{since } y \ge \alpha > \bot)
\end{aligned}$$

□

Corollary 28 *Let* $\mathsf{L} = (L, \vee, \wedge, \bot, \top)$ *be a finite lattice with some* $\alpha \in \mathsf{L}\quad, \alpha < \top$*. Then*

$$\sum_{x: x\wedge\alpha=\bot} \mu(x, \top) = 0$$

This is the *dual* to Weisner's theorem. You can prove this by standing on your head and looking at the lattice upside down:

Exercise 29 *Let* $\mathsf{L}^* := (L, \wedge, \vee, \top, \bot)$ *be the* dual *lattice of* $\mathsf{L} := (L, \vee, \wedge, \bot, \top)$*. What is* μ_{L^*} *in terms of* μ_L*? (*HINT*: Ask what is the Zeta function of* L^* *and apply a certain matrix operation to the identity* $\zeta\mu = 1$*.)*

Let us illustrate how to apply Wiesner's theorem by reworking the μ function of the lattice of sets. We shall repeatedly use the fact that isomorphic lattices have the same μ function, and that $\mu(x, y)$ only depends on the sub–interval lattice $[x, y]$.

Consider the lattice 2^U of subsets of a universal set, and let $A \subseteq B$, Observe that $[A, B]$ is isomorphic to $[\emptyset, B \setminus A]$. Hence,

$$\mu(A, B) = \mu(\emptyset, B \setminus A)$$

i.e. if we can compute $\mu(\emptyset, C)$ for a subset C we are done. For starters, we have

$$\mu(\emptyset, \emptyset) = 1$$

Now suppose $C \neq \emptyset$ then consider any $c \in C$ and take $\{c\}$ to be the special element chosen in Wiesner's Theorem. Applying it, we get:

$$\mu(\emptyset, C) + \mu(\emptyset, C \setminus \{c\}) = 0$$

Thus,

$$\mu(\emptyset, C) = -\mu(\emptyset, C \setminus \{c\}).$$

Iterating this, (or more formally, using induction) we conclude that

$$\mu(\emptyset, C) = (-1)^{|C|},$$

as we had computed before.

Finally, we shall try to compute the μ function for the lattice of partitions. This is a truly hard nut to crack, and we shall be forced to use all of the arsenal at our disposal!

First let us try to compute $\mu(\perp, \top)$ for the lattice $\mathcal{P}_n$ of partitions of $[n]$. Recall that $\top = \{12\ldots n\}$ and $\perp = \{1, 2, \ldots, n\}$. We apply Wiesner's Theorem with the special element $\Pi_0 := \{1, 23\ldots n\}$. For which partitions Π can $\Pi_0 \wedge \Pi = \perp$? An immediate answer is $\Pi = \perp$, of course. A little reflection shows that Π must come from the set $S \cup \{\perp\}$, where $S := \{\{, \{12, 3, \ldots, n\}, \{13, 2, \ldots, n\}, \ldots, \{1n, 2, \ldots, n-1\}\}$. Applying Wiesner's Theorem, we get:

$$\mu(\perp, \top) = -\sum_{x \in S} \mu(x, \top)$$

Take any $x \in S$, say the one which has the block $1i$. In any partition between x and $\top$, these two remain stuck together, hence we can identify them as a single compound entity. Then the interval lattice $[x, \top]$ is isomorphic to the partition lattice on $n-1$ elements. Hence, denoting $\mu_n(\perp, \top)$ for $\mathcal{P}_n$ by μ_n, for $n \geq 1$, we have

$$\mu_n = -(n-1)\mu_{n-1},$$

whose solution is

$$\mu_n = (-1)^{n-1}(n-1)!.$$

Consider now, two partitions Π_1, Π_2 with $\Pi_1 \leq \Pi_2$. Say that Π_2 has k blocks that when numbered are the unions of $n_1, n_2, \ldots, n_k$ blocks of Π_1. Then the interval $[\Pi_1, \Pi_2]$ is isomorphic to the product:

$$\Pi_{n_1} \times \Pi_{n_2} \times \cdots \times \Pi_{n_k}.$$

To see this, think of the blocks of Π_1 as compound entities. Using our product formula, we get

$$\mu(\Pi_1, \Pi_2) = (-1)^{|\Pi_2| - |\Pi_1|} \prod_{B \in \Pi_2} (n_B - 1)!,$$

where n_B denotes the number of blocks of Π_1 that are contained in a block B of Π_2 and recalling that for partitions, the rank function is n minus the number of blocks.

Exercise 30 *Carry out the slightly tricky verification of these statements.*

Now that we've sweated to compute this μ function, we'd better have some use for it!

Let (V, E) be an undirected graph. We can asume $V = [n]$. For a positive integer k, a map $\chi : V \mapsto [k]$ is called a **k-colouring**. If χ further satisfies the condition that for each edge

$(u,v) \in E$, $\chi(u) \neq \chi(v)$ then it is called a *proper k-colouring.* We shall always mean proper k colouring, when we say colouring below. We define $\chi_G(k) := |\{\chi \mid \chi_i \text{ is a } k\text{-colouring of} G\}|$.

Let (V,E) be an undirected graph. How many ways are there to decorate it with a (proper) k-colouring? That is, what is the size of $\chi_G(k)$?

Given an undirected graph $G := (V,E), V \sim [n]$, we define the partition lattice

$$\Pi_G := \{\Pi \mid \text{for each block } A \in \Pi, G_{|A} \text{ is connected}\}$$

We now define the sets $G(\Pi), F(\Pi)$ as follows:

$$G(\Pi) := \{\chi | \chi_A \text{ is constant on each } A \in \Pi\}$$

i.e. it is the set of all colourings that give the same colour to each member of any given block.

$$F(\Pi) := G(\Pi) \cap \{\chi | \chi_A \neq \chi_B \text{ if } a \in A, b \in B, (a,b) \in E\}$$

i.e. if A and B are two blocks with an edge of G going between them, then a colouring in $F(\Pi)$ must assign different colours to them. This is all getting very bizarre! However, note that we would be interested in $|F(\perp)|$. Set

$$g(\Pi) := |G(\Pi)| \quad \text{and} \quad f(\Pi) := |F(\Pi)|$$

we have $\chi_G(k) = f(\perp)$. In other words we have the solution to the problem if we can compute $f(\perp)$. It is easy to compute $g(\Pi) = k^{|\Pi|}$ (Since there are no restrictions we can pick any one of the k colours for each block). Observe now, that given an element (colouring) $\chi_0 \in G(\Pi)$ we can map it to an unique colouring in $F(\Pi')$ for some $\Pi' \geq_P \Pi$. We do this by merging every two blocks $A, B \in \Pi$ which conflict, i.e. if $a \in A, b \in B, (a,b) \in E$ but $\chi_A = \chi_B$. The other way round, every colouring in $F(\Pi')$ for a $\Pi' \geq_P \Pi$ corresponds to an unique colouring in $G(\Pi)$, namely the same colouring.
I.e. we have an bijection:

$$G(\Pi) \leftrightarrow \bigcup_{\Pi' \geq_P \Pi} F(\Pi')$$

and so we get

$$g(\Pi) = \sum_{\Pi' \geq_P \Pi} f(\Pi').$$

By Möbius inversion we have that $f = \mu g$. Remembering $\chi_G(k) = f(\perp)$ this finally gives us

$$\begin{aligned}
\chi_G(k) &= \sum_{\Pi \geq \perp} \mu(\perp, \Pi) g(\Pi) \\
&= \sum_{\Pi} \mu(\perp, \Pi) k^{|\Pi|} \\
&= \sum_{1 \leq j \leq n} \Big(\sum_{|\Pi| = j} \mu(\perp, \Pi) \Big) k^j
\end{aligned}$$

So in fact, the number of colourings is polynomial in the number of colours. This polynomial is called the *chromatic polynomial* of the graph.

Chapter 5

Generating functions

You've just ordered a pizza to be delivered and while you're waiting you start wondering in how many ways you can pay for it using only coins of denomination Dkr. 1,5,10 and 20. You know that the pizza costs Dkr 40 [1]. Letting C_i denote the amounts you can pay with coins of denomination i and smaller we can write the equations (here $\odot$ = Dkr. 1, $\bigodot$ = Dkr. 5, $\circ$ = Dkr. 10 and $\bigcirc$ = Dkr.20):

$$
\begin{aligned}
C_1 &= (\not\odot + \odot + \odot\odot + \odot\odot\odot + \ldots) \\
&= (\not\odot + \odot + \odot^2 + \odot^3 \ldots) \\
C_5 &= (\not\bigodot + \bigodot + \bigodot^2 + \ldots)C_1 \\
C_{10} &= (\not\circ + \circ + \circ^2 + \ldots)C_5 \\
C_{20} &= (\not\bigcirc + \bigcirc + \bigcirc^2 + \ldots)C_{10}
\end{aligned}
$$

and count the number of coin combinations in C_{20} yielding 40. This is clearly the number of ways to pay 40 kr using only coins of denomination Dkr 1,5,10 and 20. The problem is, you might run out of paper trying to compute C_{20}! Solution: throw in some algebra to do the work for you!

Letting z denote Dkr. 1, z^5, z^{10} and z^{20} represent respectively Dkr. 5, 10 and 20, we can write the equations above as functions where the coefficient of z^n in C_i tells us in how many ways we can pay for the amount Dkr. n using coins of denomination i or smaller. So, our equations now become:

$$
\begin{aligned}
C_1(z) &= (1 + z + z^1 + z^2 + \ldots) \\
C_5(z) &= (1 + z^5 + z^{10} + \ldots)C_1(z) \\
C_{10}(z) &= (1 + z^{10} + z^{20} + \ldots)C_5(z) \\
C_{20}(z) &= (1 + z^{20} + z^{40} + \ldots)C_{10}(z) \qquad (5.1)
\end{aligned}
$$

We are interested in the coefficient of z^{40} in $C_{20}(z)$.

Given a sequence of numbers:

$$a_0, a_1, a_2, \ldots, a_n, \ldots$$

[1] Poor budget "Turbo–Sprint" pizza from down the road!

their *generating function* is a power series in one variable which looks like this:

$$A(z) := a_0 + a_1 z + a_2 z^2 + ...$$

The number a_n in the original sequence is the coefficient of z^n in $A(z)$. To extract the original numbers out of the generating function, we introduce the operator:

$$< z^n > A(z) := a_n$$

which pick out the n'th coefficient.

In our example, the generating function for the number of ways to make change using kr. 1, 5, 10 and 20 is $C_{20}(z)$. And our answer is:

$$< z^{40} > C_{20}(z) \qquad := \qquad C_{20,40}$$

But still we have to compute it and to do that let us rewrite (5.1) as:

$$\begin{aligned}
(1-z)C_1(z) &= 1 \\
(1-z^5)C_5(z) &= C_1(z) \\
(1-z^{10})C_{10}(z) &= C_5(z) \\
(1-z^{20})C_{20}(z) &= C_{10}(z)
\end{aligned}$$

If we have two generating functions $F(z)$ and $G(z)$ such that $F(z) = G(z)$ then $< z^n > F(z) =< z^n > G(z)$. Moreover, the operator $< z^n >$ is *linear* in the following way:

$$< z^n > (aF(z) + bG(z)) = a < z^n > F(z) + b < z^n > G(z).$$

Also multiplying by z corresponds to a *shift*:

$$< z^n > zF(z) =< z^{n-1} > F(z). \tag{5.2}$$

The reason for the calling it a shift is that if $A(z) = a_0 + a_1 z + a_2 z^2 + \cdots$ is the generating function for the sequence $a_0, a_1, a_2, \ldots$, then $zA(z)$ is the generating function for the shifted sequence $0, a_0, a_1, a_2, \ldots$ as the previous formula shows.

Applying the operator $< z^n >$ to both sides of the previous set of change equations, and observing that $< z^n > 1 = [n = 0]$, we obtain (here we are denoting the sequence generated by the function $C_i(z)$, in the natural way as $C_{i,0}, C_{i,1}, C_{i,2}, \ldots$:

$$\begin{aligned}
C_{1,n} &= C_{1,n-1} + [n = 0] \\
C_{5,n} &= C_{5,n-5} + C_{1,n} \\
C_{10,n} &= C_{10,n-10} + C_{5,n} \\
C_{20,n} &= C_{20,n-20} + C_{10,n}
\end{aligned}$$

To compute $C_{20,40}$ we can unwind this set of equations recursively:

$$\begin{aligned}
C_{20,40} &= C_{20,20} + C_{10,40} \\
&= C_{20,0} + C_{10,20} + C_{10,30} + C_{5,40} \\
&= C_{20,0} + C_{10,10} + C_{5,20} + C_{10,20} + C_{5,30} + C_{5,35} + C_{1,40} \\
&\vdots
\end{aligned}$$

The values needed to be computed to find $C_{20,40}$ are computed in the diagram below and we discover that there are 35 ways to pay the pizza man using coins only of denomination kr 1,5,10 and 20.

$\backslash n$	0	5	10	15	20	25	30	35	40
C_1	1	1	1	1	1	1	1	1	1
C_5	1	2	3	4	5	6	7	8	9
C_{10}	1		4		9		16		25
C_{20}	1				10				35

What is the answer in some other country, say India? Given some coins of denominations $a_1 < a_2 < a_3 < \cdots$ how many ways are there to make change for a given amount? Can we make change for every amount?

Theorem 31 (Schur's Theorem) *If the numbers $(a_1, a_2, a_3, ...)$ are relatively prime, i.e. $gcd(a_1, a_2, a_3, \ldots) = 1$ then for any sufficiently large n one can make change using these coins.*

To illustrate this let us take an example.

Example 32 *Let $a = 2$ and $b = 5$ then indeed $gcd(2, 5) = 1$ and we can make change for any sufficiently large amount. But note that we cannot make change for the amount 3.*

Of course the theorem did say for n large and most people would agree that 3 is not very large! So what exactly *is* "sufficiently large"? The answer in general is unknown, but in the case that there are only two denominations, as in our example, there is a simple answer:

Theorem 33 *Let $a < b$ $gcd(a, b) = 1$ then:*

- *For any $n \geq (a-1)(b-1)$ one can make change.*
- *$(a-1)(b-1) - 1$ cannot be changed.*
- *Of the numbers less than or equal to $(a-1)(b-1) - 1$ exactly half can be changed.*

Let us now move into a fantasy world where we can have coins of *every* denomination [2]. This brings us into the fascinating realm of *the theory of partitions*, a field in which generating functions have achieved some of their most celebrated triumphs.

A *partition of n* is a way to write n as a sum of positive integers:

$$n = a_1 + a_2 + ... + a_k, \quad a_1 \geq a_2 \geq a_3 \geq ... \geq 1.$$

We denote by $p(n)$, the number of partitions of n.

Example 34 *Partitions of 5.*

$$\begin{aligned} 5 &= 5 \\ &= 4+1 \\ &= 3+2 \end{aligned}$$

[2] Of course, we are already in a fantasy world, assuming as we always have so far, an unlimited supply of coins in our pockets!

$$\begin{aligned}
&= 2+2+1 \\
&= 3+1+1 \\
&= 2+1+1+1 \\
&= 1+1+1+1+1
\end{aligned}$$

That is:

$$p(5) = 7$$

To find the number of partitions of n as a function of n we first realize that the numbers we are interested in are

$$p(1)\ p(2)\ p(3) \cdots p(n) \cdots$$

We immediately introduce the corresponding generating function $P(z)$ for this sequence:

$$P(z) = \sum_{n \geq 0} p(n) z^n$$

But how do we compute $P(z)$? The intuition is exactly the same as for the pizza example only here we have coins of every denomination to pay for the pizza costing n (in whatever the currency is)!

$$\begin{aligned}
P(z) &= (\cancel{1} + 1 + 11 + 111 + ...)(\cancel{2} + 2 + 22 + 222 + ...)...(\cancel{n} + n + nn + nnn + ...)... \\
&= (1 + z + z^2 + z^3 + ...)(1 + z^2 + z^4 + z^6 + ...)...(1 + z^n + z^{2n} + ...)... \\
&= (\frac{1}{1-z})(\frac{1}{1-z^2})...(\frac{1}{1-z^n})... \\
&= \prod_{m \geq 1} \frac{1}{1-z^m}
\end{aligned}$$

Example 35 *The odd partitions of 6. How many partitions, $p_o(6)$ of 6 are there into* odd *parts? We can list them out:*

$$\begin{aligned}
6 &= 5+1 \\
&= 3+3 \\
&= 3+1+1+1 \\
&= 1+1+1+1+1+1
\end{aligned}$$

And we see that $p_o(6) = 4$.

In general, how many partitions $p_o(n)$ of n are there, into odd parts? The generating function for this problem is very much like the one for the partition problem except that here we are only allowed to use the odd numbered coins. Thus the generating function is:

$$\begin{aligned}
P_o(z) &= (\cancel{1} + 1 + 11 + 111 + ...)(\cancel{3} + 3 + 33 + 333 + ...)... \\
&= (1 + z + z^2 + z^3 + ...)(1 + z^3 + z^6 + ...)... \\
&= (\frac{1}{1-z})(\frac{1}{1-z^3}) \cdots \\
&= \prod_{\substack{m \\ \text{odd}}} \frac{1}{1-z^m}
\end{aligned}$$

Example 36 *The partitions of 6 into unequal parts. How many partitions of* 6 *are there there into unequal parts? Denoting this by* $p_u(6)$*, we can find it by listing out these partitions:*

$$\begin{aligned} 6 &= 6 \\ &= 5+1 \\ &= 4+2 \\ &= 3+2+1 \end{aligned}$$

And we see that $p_u(6) = 4$*.*

To find the generating function we are again inspired by the coin problem, but here we may use each coin denomination *at most once* i.e. once or not at all.

$$\begin{aligned} P_u(z) &= (\cancel{1}+1)(\cancel{2}+2)(\cancel{3}+3)... \\ &= (1+z)(1+z^2)(1+z^3)... \\ &= \prod_{m\geq 0}(1+z^m) \end{aligned}$$

The two examples show that $p_o(6) = p_u(6) = 4$ but does this hold in general? Let's look at the generating functions again. We start with $P_o(z)$, multiply it in the numerator and denominator by the missing even power terms, factor $(1-z^{2m}) = (1-z^m)(1+z^m)$, and finally use the wonderful properties of infinite products:

$$\begin{aligned} P_o(z) &= (\frac{1}{(1-z)})(\frac{1}{(1-z^3)})(\frac{1}{(1-z^5)})\cdots \\ &= \frac{(1-z^2)(1-z^4)(1-z^6)\cdots}{(1-z)(1-z^2)(1-z^3)(1-z^4)(1-z^5)(1-z^6)\cdots} \\ &= \frac{(1-z)(1+z)(1-z^2)(1+z^2)(1-z^3)(1+z^3)\cdots}{(1-z)(1-z^2)(1-z^3)(1-z^4)(1-z^5)(1-z^6)\cdots} \\ &= (1+z)(1+z^2)(1+z^3)... \\ &= P_u(z) \end{aligned}$$

Amazingly, $P_o(z) = P_u(z)$, although they didn't at all look the same to start with! Of course this means that $p_o(n) = p_u(n)$ i.e. *the number of partitions of a number into odd parts equals the number of its partitions inot parts all of which are unequal* !

Euler was a master at proving such amazing results. A very interesting result of his is:

$$\prod_{m\geq 0}(1-z^m) = 1 + \sum_{k\geq 1}(-1)^k(z^{\omega(k)} + z^{\omega(-k)}) \tag{5.3}$$

Here $\omega(k) := (3k^2+k)/2$ and $\omega(-k) := (3k^2-k)/2$ are the so–called *pentagonal numbers*. The left hand side here looks familiar; indeed it is the reciprocal of the generating function for partitions! Multiplying them together means that

$$\sum_n p(n)z^n \cdot (1 + \sum_{k\geq 1}(-1)^k(z^{\omega(k)} + z^{\omega(-k)})) = 1 \tag{5.4}$$

The stage is ripe to apply the operator $< z^n >$ to both sides of this equation.

To do this, we shall show a general rule about how the operator $< z^n >$ behaves with respect to a product. If

$$A(z) := a_0 + a_1 z + a_2 z^2 + \cdots = \sum_n a_n z^n,$$

and

$$B(z) := b_0 + b_1 z + b_2 z^2 + \cdots = \sum_n b_n z^n,$$

then by carrying out the multiplication as we learnt in school,

$$A(z)B(z) = a_0 b_0 + (a_0 b_1 + a_1 b_0) z^2 + \cdots = \sum_n (\sum_{0 \le k \le n} a_k b_{n-k}) z^n.$$

That is,

$$< z^n > A(z)B(z) = \sum_k a_k b_{n-k}.$$

The sequence whose nth term is $\sum_k a_k b_{n-k}$ is called the *convolution* of the two original sequences of the a's and b's. Thus convoluting sequences corresponds to simple multiplication of generating functions!

Applying this observation to (5.4), we get:

$$\sum_k (-1)^k (p(n - \omega(k)) + p(n - \omega(-k))) = 0.$$

This actually gives a recursion formula for p_n:

$$p(n) = p(n-1) + p(n-2) - p(n-5) - p(n-7) + \cdots$$

and we can use it to rapidly compute the values of p_n:

n	0	1	2	3	4	5	6	7	8	9
$p(n)$	1	1	2	3	5	7	11	15	22	30

Let us now go on a small guided tour trying to get to know generating functions and how to manipulate them. We shall record our observations in a travel diary which is Table 5.

Lets start with the simplest (non–trivial) sequence of all: a 1 followed by zeroes. The generating function is just 1. This is our first entry. Not very exciting, but have patience! Let's divide this by $(1 - z)$. What sequence do we get? Well, we know by summing geometric series that

$$\frac{1}{1-z} = 1 + z + z^2 + \cdots = \sum_n z^n,$$

That is, the sequence generated is all ones. This is our second entry.

Notice, that we can view the sequence of all ones in another way: take the previous sequence (1 followed by zeroes) and form the *cumulative* sum upto any point. In fact this works in general: Given a generating function $A(z)$ for a sequence of numbers

$$A(z) = a_0 + a_1 z + a_2 z^2 + \cdots = \sum_{n \ge 0} a_n z^n$$

what sequence does $B(z) := A(z)/(1-z)$ generate? Well, remembering the rule that products of generating functions correspond to convolutions of the corresponding sequences, we see that

$$< z^n > B(z) = \sum_k a_{n-k} = a_0 + a_1 + \cdots + a_n.$$

We can try repeat the trick; dividing our second entry by $(1-z)$ again yields the generating function for the sequence $1, 2, 3, \cdots$, whose nth term is $(n+1)$. To normalise this, we multiply by z and get the generating function for the natural numbers: $z/(1-z)^2$, our third entry. Things are beginning to get more interesting.

Let's do it once more! Divide by $(1-z)$ and we get the generating function for the partial sums $0, 0+1, 0+1+2, \ldots$:

$$\sum_n (\sum_{k \leq n} k) z^n = z/(1-z)^3.$$

Our third entry. Well, what's the nth term here?

$$\begin{aligned} < z^n > \frac{z}{(1-z)^3} &= < z^{n-1} > \frac{1}{(1-z)^3} \\ &= < z^{n-1} > (1-z)^{-3} \end{aligned}$$

But how can we get further? The Binomial theorem helps us.

Theorem 37 (Binomial Theorem) *For any a, b and $n > 0$,*

$$(a+b)^n = \sum_k \binom{n}{k} a^{n-k} b^k.$$

As important special cases, we have

$$(1+x)^n = \sum_k \binom{n}{k} x^k,$$

and

$$(1-x)^n = \sum_k \binom{n}{k} (-1)^k x^k.$$

Here of course, for any $n \geq 0$

$$\begin{aligned} \binom{n}{k} &:= \frac{n!}{k!(n-k)!} \\ &= \frac{n \cdot (n-1) \cdots (n-k+1)}{k!} \end{aligned}$$

The last expression allows us to make sense of $\binom{r}{k}$ even when r is not a positive integer. This yields an extended Binomial Theorem valid even when r is not a positive integer. As a special case, it can be verified that if we have a top index which is a negative integer, then we can use:

$$\binom{-n}{k} = (-1)^k \binom{k+n-1}{k}.$$

Exercise 38 *Verify this transformation.*

Now we are back in business and can compute:

$$\begin{aligned} < z^{n-1} > (1-z)^{-3} &= < z^{n-1} > \sum_k (-1)^k \binom{-3}{k} z^k \\ &= (-1)^{n-1} \binom{-3}{n-1} \\ &= (-1)^{n-1}(-1)^{n-1} \binom{n-1+3-1}{n-1} \\ &= \binom{n+1}{n-1} \end{aligned}$$

If this looks a bit unfamiliar, perhaps you're happier with the answer

$$\binom{n+1}{n-1} = n(n+1)/2.$$

Sequence	*Generating Function*
$1, 0, 0, 0, \ldots$	1
$1, 1, 1, 1, \ldots$	$\frac{1}{1-z}$
$0, 1, 2, 3, \ldots$	$\frac{z}{(1-z)^2}$
$0, 0+1, 0+1+2, 0+1+2+3, \ldots$	$\frac{z}{(1-z)^3}$

Exercise 39 *How many solutions in non–negative integers are there to the equation*

$$x_1 + x_2 + \cdots + x_k = n,$$

where k, n are positive integers?

Chapter 6

Generating functions II

In the previous chapter we had embarked on our tour of "Generatingfunctionland". Here, we continue on our way.

Let's do the "differentiate, multiply and divide" trick once again. Start with our by–now–familiar friend:

$$\sum_{k\geq 0} z^k = \frac{1}{1-z},$$

differentiate and multiply by z:

$$\sum_{k\geq 0} k \cdot z^k = \frac{z}{(1-z)^2}.$$

Once again:

$$\sum_{k\geq 0} k^2 \cdot z^k = \frac{z\cdot(1+z)}{(1-z)^3}$$

The sequence here is $0^2, 1^2, 2^2, 3^2, \ldots$, the squares of the natural numbers. To form the cumulative sums, we divide by $1-z$:

$$\sum_{n\geq 0}(\sum_{k\leq n} k^2)z^n = \frac{z\cdot(1+z)}{(1-z)^4}.$$

Let's pick out the nth term using the $< z^n >$ operator. Remember the linearity and shift proerties of $< z^n >$?

$$\begin{aligned} < z^n > \frac{z\cdot(1+z)}{(1-z)^4} &= < z^n > \frac{z}{(1-z)^4} + < z^n > \frac{z^2}{(1-z)^4} \\ &= < z^{n-1} > (1-z)^{-4} + < z^{n-2} > (1-z)^{-4}. \end{aligned}$$

Now use the Binomial Theorem to see that the coefficient we want is:

$$\binom{-4}{n-1}\cdot(-1)^{n-1} + \binom{-4}{n-2}(-1)^{n-2}.$$

We don't like negative values at the top of our binomial coefficients, so we convert it using the transformation derived in exercise 38:

$$\binom{n+2}{n-1} + \binom{n+1}{n-2}$$

Which by expansion gives us the familiar answer:

$$1^2 + 2^2 + 3^2 + \cdots + n^2 = \frac{1}{6} \cdot n \cdot (n+1) \cdot (2 \cdot n + 1).$$

If we used differentiation, can integration be far behind? Our starting point is again:

$$\sum_{k \geq 0} z^k = \frac{1}{1-z}.$$

Integrating (or anti–differentiating) this,

$$\int \sum_{k \geq 0} z^k = \int \frac{1}{1-z},$$

i.e.

$$\sum_{k \geq 0} \frac{z^{k+1}}{k+1} = \ln \frac{1}{1-z}.$$

or, rewriting the dummy index,

$$\sum_{k \geq 1} \frac{z^k}{k} = \ln \frac{1}{1-z}.$$

Forming the cumulative partial sums as usual,

$$\sum_{n \geq 1} \Big(\sum_{k \leq n} \frac{1}{k} \Big) z^n = \frac{1}{1-z} \ln\Big(\frac{1}{1-z}\Big).$$

The coefficients in this series,

$$\frac{1}{1-z} \ln\Big(\frac{1}{1-z}\Big) := H_1 z + H_2 z^2 + H_3 z^3 + \cdots + H_n z^n + \cdots,$$

namely,

$$H_n := 1 + 1/2 + 1/3 + \cdots + 1/n,$$

are called the *harmonic numbers.*

Let's take the example that was used to advertise this course, one of the all time classics! The *Fibonacci numbers* are given by the recurrence

$$f_n = f_{n-1} + f_{n-2} \quad \text{for } n \geq 2, \tag{6.1}$$

where $f_0 = 0$ and $f_1 = 1$. We plunge in and introduce the generating function straight away:

$$F(z) = \sum_{n \geq 0} f_n \cdot z^n.$$

We convert (6.1) which relates the numbers into a relation involving the generating function instead. Multiply both sides by z^n and sum over the allowed range, namely, $n \geq 2$:

$$\sum_{n \geq 2} f_n \cdot z^n = \sum_{n \geq 2} f_{n-1} \cdot z^n + \sum_{n \geq 2} f_{n-2} \cdot z^n. \tag{6.2}$$

Let's look at what the terms are on both sides of this equation. First,

$$\begin{aligned}\sum_{n\geq 2} f_n \cdot z^n &= \sum_{n\geq 0} f_n \cdot z^n - f_0 - f_1 z \\ &= F(z) - z. \qquad (6.3)\end{aligned}$$

Similarly,

$$\begin{aligned}\sum_{n\geq 2} f_{n-1} z^n &= z \sum_{n\geq 2} f_{n-1} z^{n-1} \\ &= z \sum_{n\geq 1} f_n z^n \\ &= z(F(z) - f_0) \\ &= zF(z). \qquad (6.4)\end{aligned}$$

and finally,

$$\begin{aligned}\sum_{n\geq 2} f_{n-2} z^n &= z^2 \sum_{n\geq 2} f_{n-2} z^{n-2} \\ &= z^2 \sum_{n\geq 0} f_n z^n \\ &= z^2 F(z) \qquad (6.5)\end{aligned}$$

Substituting (6.3),(6.4) and (6.5) into (6.2), we get an equation satisfied by $F(z)$:

$$F(z) - z = zF(z) + z^2 F(z).$$

Solving this equation, we get

$$F(z) = \frac{z}{1 - z - z^2}.$$

We are interested in the nth number:

$$\begin{aligned}< z^n > F(z) &= < z^n > \frac{z}{1 - z - z^2} \\ &= < z^{n-1} > \frac{1}{1 - z - z^2}.\end{aligned}$$

Unfortunately we can't apply the binomial theorem here so instead we use *partial fractions*. First we look at the quadratic equation:

$$z^2 + z - 1 = 0,$$

whose solution is:

$$z = -\phi, -\hat{\phi},$$

where

$$\phi := \frac{1 + \sqrt{5}}{2}, \hat{\phi} := \frac{1 - \sqrt{5}}{2}.$$

The number ϕ held special importance for the architects of antiquity and is called the *golden ratio*. Note that

$$\phi + \hat{\phi} = 1 = -\phi\hat{\phi}.$$

Hence we can factor

$$z^2 + z - 1 = (z + \phi) \cdot (z + \hat{\phi})$$

For the reciprocal, by multiplying and dividing by $\phi\hat{\phi} = -1$, we have:

$$\begin{aligned} \frac{1}{1 - z - z^2} &= \frac{-1}{(z + \phi)(z + \hat{\phi})} \\ &= \frac{1}{(\phi z - 1)(\hat{\phi} z - 1)} \\ &= \frac{1}{(1 - \phi z)(1 - \hat{\phi} z)} \end{aligned}$$

Now we *decompose* into partial fractions:

$$\frac{1}{(1 - \phi z)(1 - \hat{\phi} z)} = \frac{A}{1 - \phi z} + \frac{\hat{A}}{1 - \hat{\phi} z}.$$

which can only be possible if the yet undetermined constants $A, \hat{A}$ satisfy

$$A(1 - \hat{\phi} z) + \hat{A}(1 - \phi z) = 1.$$

This in turn is only possible if

$$A + \hat{A} = 1 \quad \text{and} \quad \hat{\phi} A + \phi \hat{A} = 0.$$

The solution to this pair of linear equations is:

$$A := \frac{\phi}{\phi - \hat{\phi}} \quad \text{and} \quad \hat{A} := \frac{\hat{\phi}}{\hat{\phi} - \phi}.$$

Well back to where we came from, now we can find the nth Fibonacci number:

$$\begin{aligned} < z^n > \cdot F(z) &= < z^{n-1} > \frac{1}{1 - z - z^2} \\ &= < z^{n-1} > \frac{1}{\sqrt{5}} \left(\frac{\phi}{1 - \phi \cdot z} - \frac{\hat{\phi}}{1 - \phi \cdot z} \right) \\ &= \frac{1}{\sqrt{5}} \left(< z^{n-1} > \frac{\phi}{1 - \phi \cdot z} - < z^{n-1} > \frac{\hat{\phi}}{1 - \hat{\phi} \cdot z} \right) \\ &= \frac{1}{\sqrt{5}} (\phi^n - \hat{\phi}^n) \end{aligned}$$

In the last line we used the Binomial Theorem once again. This might seem a silly answer because the Fibonacci numbers are natural numbers. In fact this value is always an integer, and the next exercise asks you to check this by deriving another form of the answer.

Exercise 40 *Show that f_n is the nearest integer to $\frac{1}{\sqrt{5}} \phi^n$, that is,*

$$f_n = \lfloor \frac{1}{\sqrt{5}} \cdot \phi^n + \frac{1}{2} \rfloor.$$

In computer science, a common problem is that of balancing parentheses. For example, when typesetting this document in TeX, the editor complains whenever it detects that you have mis–balanced parentheses! Another example is in Mathematics: suppose we have a binary operation $\otimes$ that does not obey the associative law. Then for an expression $(m_0 \otimes m_1 \otimes ... \otimes m_n)$ to make sense, it must be properly parenthesised. In how many different ways can we do this, leading to potentially all different results? Let c_n denote this number, and introduce right away, the corresponding generating function $C(z) := \sum_n c_n \cdot z^n$

We can easily compute c_n for small numbers n:

$$c_0 = 1 : (m_0),$$

$$c_1 = 1 : (m_0 \otimes m_1),$$

$$c_2 = 2 : (m_0 \otimes m_1) \otimes m_2, m_0 \otimes (m_1 \otimes m_2).$$

We notice that any legal parenthesising must have a final $\otimes$ operation, the place which this last $\otimes$ operation is located we call k.

$$((\underbrace{m_0 ... m_k}_{k+1}) \otimes (\underbrace{m_{k+1} ... m_n}_{n-k})))$$

as k can be in the range of $0 \leq k < n$ we have the following recurrence:

$$c_n = \sum_k c_k \cdot c_{n-k-1}, \quad \text{for } n > 0.$$

We convert this relation between numbers into a relation between generating functions as before.

$$\sum_{n>0} c_n \cdot z^n = \sum_{n>0} (\sum_k c_k \cdot c_{n-k-1}) \cdot z^n.$$

The left hand side of this equation is simply $C(z) - 1$ since $c_0 = 1$. For the right side, we notice that $\sum_k c_k \cdot c_{n-k-1}$ looks very familiar: it is the convolution of the sequence c_n with itself! More precisely, our rules for convolutions of sequences and the corresponding generating functions tells us that

$$\sum_k c_k \cdot c_{n-k-1} =< z^{n-1} > (C(z))^2.$$

Hence, our relation between generating functions is:

$$\begin{aligned} C(z) - 1 &= \sum_{n>0} (< z^{n-1} > (C(z))^2) z^n \\ &= z \sum_{n>0} (< z^{n-1} > (C(z))^2) z^{n-1} \\ &= z \sum_{n \geq 0} (< z^n > (C(z))^2) z^n \\ &= z(C(z))^2 \end{aligned}$$

In the last line, we decipher the notation in

$$A(z) = \sum_n (< z^n > A(z)) z^n.$$

So we have a quadratic equation again:

$$z \cdot C(z)^2 - C(z) + 1 = 0$$

But wait a minute! The unknown here is a *generating function series* not a number! Never mind – we are going to be absolutely democratic and treat all as equal! The quadratic formula says that

$$C_+(z) = \frac{1 + \sqrt{1 - 4 \cdot z}}{2 \cdot z},$$

or

$$C_-(z) = \frac{1 - \sqrt{1 - 4 \cdot z}}{2 \cdot z}.$$

Which one is it? Normally, one would choose the positive sign. But here this is a disaster. For $z = 0$ this function is unbounded, whereas it should give $C(0) = c_0$. So we are forced to take refuge in the other sign. Does it help? At first sight, this is bad too, but at $z = 0$ it's of the form $\frac{0}{0}$, and there's some hope! In fact, a dose of L'Hôpital's rule shows that $\lim_{z \to 0} C_-(z) = 1$ and we are saved! Our generating function is

$$C(z) = \frac{1 - \sqrt{1 - 4 \cdot z}}{2 \cdot z}.$$

All that remains is to extract c_n out of this. We shall use the Binomial Theorem to do this, but this time it will take a bit more work!

$$\begin{aligned}
< z^n > C(z) &= < z^n > \frac{1 - (1 - 4 \cdot z)^{1/2}}{2 \cdot z} \\
&= < z^n > \frac{1 - \sum_k \binom{1/2}{k}(-4z)^k}{2 \cdot z} \\
&= < z^n > \frac{\sum_{k>0} \binom{1/2}{k}(-4z)^k}{2 \cdot z} \\
&= -\frac{1}{2} < z^{n+1} > \sum_{k>0} \binom{1/2}{k}(-4z)^k \\
&= -\frac{1}{2} \cdot \binom{1/2}{n+1}(-4)^{n+1}.
\end{aligned}$$

Now we seem to be in deep shit! The $\binom{1/2}{n+1}$ is real mean and ugly looking! But all we can do is to resolutely press ahead. Remembering how the binomial coefficient is defined for non–integral upper parameters, we have:

$$\begin{aligned}
\binom{1/2}{n+1} &= \frac{1/2 \cdot (1/2 - 1) \cdot (1/2 - 2) \cdot \cdots \cdot (1/2 - (n+1))}{(n+1)!} \\
&= \frac{1}{2^{n+1}} \cdot \frac{1 \cdot (-1) \cdot (1 - 4) \cdots (1 - 2 \cdot n)}{(n+1)!} \\
&= \frac{1}{2^{n+1}} \cdot \frac{1 \cdot (-1) \cdot (-3) ... (-(2n-1))}{(n+1)!} \\
&= \frac{(-1)^n}{2^{n+1}} \cdot \frac{1 \cdot 3 \cdot 5 \cdot ... \cdot (2n-1)}{(n+1)!}
\end{aligned}$$

The numerator has all the odd numbers, so we throw in all the even ones as well:

$$\begin{aligned}
\binom{1/2}{n+1} &= \frac{(-1)^n}{2^{n+1}} \cdot \frac{1 \cdot 2 \cdot 3 \cdot \dots \cdot (2n)}{2 \cdot 4 \cdot 6 \cdot (2 \cdot n)(n+1)!} \\
&= \frac{(-1)^n}{2 \cdot 4^n} \cdot \frac{(2n)!}{n! \cdot n+1!} \\
&= \frac{(-1)^n}{2 \cdot 4^n} \cdot \frac{(2n)!}{n! \cdot n!} \cdot \frac{1}{n+1} \\
&= \frac{(-1)^n}{2 \cdot 4^n} \cdot \binom{2n}{n} \cdot \frac{1}{n+1}
\end{aligned}$$

This is just what we seek, as things neatly fit together when we substitute back and get:

$$\begin{aligned}
< z^n > C(z) &= -\frac{1}{2} \cdot \frac{(-1)^n}{2 \cdot 4^n} \cdot \binom{2n}{n} \cdot \frac{1}{n+1} \cdot (-4)^{n+1} \\
&= \binom{2 \cdot n}{n} \cdot \frac{1}{n+1}
\end{aligned}$$

So we finally have our answer to the problem of balanced parenthesis:

$$c_n = \frac{1}{n+1} \cdot \binom{2 \cdot n}{n}$$

This is another famous sequence of numbers: the *Catalan numbers*. They crop up in other problems as well.

Exercise 41 *Show that c_n is also the number of*

- *binary trees with exactly n leaves (each node in the tree has 0 or 2 children); and*
- *the number of ways to dissect a n–gon with non–intersecting diagonals.*

Chapter 7

Yet more on Generating functions

Last time, we saw that the number of ways to legally paranthesise an expression with $n+1$ elements was equal to the nth Catalan number. We employed generating functions to do this. This time, we will take a geometrical view of the same problem.

Consider the expression:

$$m_0 \otimes m_1 \otimes \ldots \otimes m_n.$$

In how many ways can one place parenthesis in this expression? To get a visual impression of the problem consider the graph below. We shall represent a possible way to parenthesise the expression by a path in the $x-y$ grid. The horizontal axis corresponds to "time", i.e. the number of the parenthesis, and the vertical direction to the type of parenthesis (left or right). A step upwards corresponds to a left parenthesis, and a step downward to a right parenthesis. That is, a line segment from $(x, y) \to (x+1, y+1)$ corresponds to a left-bracket while a segment from $(x, y) \to (x+1, y-1)$ corresponds to a right-bracket. Thus the path is made up of unit line–segments of slope $+1$ or -1. Now clearly the number of ways to put the brackets (legally or otherwise) is equal to the number of paths from $(0, 0)$ to $(2n, 0)$, because we must use n pairs of brackets overall. Now, we observe crucially that *the legal parenthesisings correspond to paths that do not cross the x-axis.* Thus we must count the number of paths from $(0, 0)$ to $(2n, 0)$ that do not cross the x–axis.

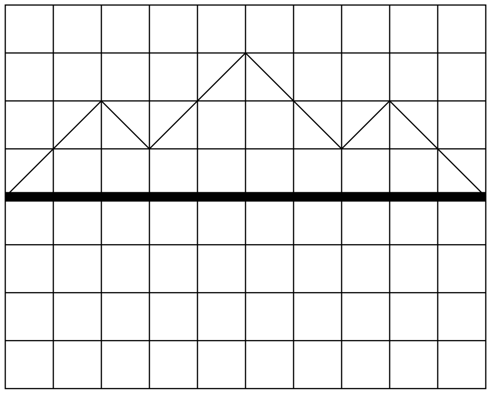

a legal paranthesing path

First we observe:

$$\#\{\text{paths from}(0,0) \to (2n,0)\} = \binom{2n}{n}, \tag{7.1}$$

because this is the same as the number of ways to interleave n left brackets with n right brackets in all possible ways. In general, let us try to extend this to computing the number of all paths $(x_1, y_1) \to (x_2, y_2)$, for $x_2 > x_1$. If u is the number of "up" steps and d the number of "down" steps in such a path, then of course, we have

$$u + d = x_2 - x_1 \quad \text{and} \quad y_1 + u - d = y_2. \tag{7.2}$$

That is, such a path exists iff there are non–negative integers u, d satisfying (7.2). And when they do, they are uniquely determined by solving the set of linear equations to get:

$$u = \frac{(x_2 - x_1) + (y_2 - y_1)}{2}, d = \frac{(x_2 - x_1) - (y_2 - y_1)}{2}. \tag{7.3}$$

Thus, we get:

$$\#\{\text{paths from}(x_1, y_1) \to (x_2, y_2), \} = \begin{cases} \binom{u+d}{u}, & \text{if } x_2 > x_1; \\ 0 & \text{otherwise.} \end{cases} \tag{7.4}$$

Returning to the original problem we are interested in finding the number of "bad" paths out of the $\binom{2n}{n}$ paths in all. First notice that any path is "bad", if it crosses the x axis. Let $(x, -1)$ be the first such point on the path. We now perpetrate a clever trick.

We *reflect* the initial portion $((0,0) \to (x, -1))$ of the path, in the line $y = -1$. In this way, observe that we get a *one-to-one correspondence between paths which cross the x–axis and the set of all paths from* $(0, -2)$ *to* $(2n, 0)$. Why is this correspondence one–one? Consider two different paths that cross the x–axis. If they cross over for the first time at different points, then the corresponding reflected points do too and hence are different. If they cross at the same point, they either differ on the rest of the path or the initial portion before the crossover. In the former case, the corresponding reflected paths are different as well, as they agree with the original paths on the remaining portion. Finally, in the latter case, the reflected portions are faithful mirror images of the original paths, hence they differ if the originals do. Thus,

$$\begin{aligned} \#\{\text{bad paths}\} &= \#\{\text{paths from } (0, -2) \to (2n, 0)\} \\ &= \binom{2n}{n+1}, \quad \text{using (7.4).} \end{aligned}$$

Finally, then,

$$\begin{aligned} \#\{\text{good paths}\} &= \binom{2n}{n} - \binom{2n}{n+1} \\ &= \frac{2n!}{n!n!} - \frac{2n!}{(n+1)!(n-1)!} \\ &= \frac{2n!}{n!n!}(1 - \frac{n}{n+1}) \\ &= \binom{2n}{n}\frac{1}{n+1}. \end{aligned}$$

which is what we got last time! This neat little combinatorial argument seems to make a mockery of our gymnastics with generating functions last time. But first, the approach via generating functions is systematic whereas this one hinged on a clever trick (the reflection).

Second, we shall see that the information on generating functions can be useful for other things as well.

Generating functions can be used to prove combinatorial identities. For starters, let's try to find:

$$\sum_k \binom{k}{n-k}$$

Let us denote $f(n) := \sum_k \binom{k}{n-k}$ and mechanically write down the generating function $F(z) := \sum_n f(n)z^n$. Now,

$$\begin{aligned}
F(z) &= \sum_n f(n)z^n \\
&= \sum_n \sum_k \binom{k}{n-k} z^n \\
& \quad \text{interchange the order of summation} \\
&= \sum_k \sum_n \binom{k}{n-k} z^n \\
&= \sum_k z^k \sum_n \binom{k}{n-k} z^{n-k} \\
&= \sum_k z^k \sum_r \binom{k}{r} z^r \\
& \quad \text{Binomial Theorem} \\
&= \sum_k z^k (1+z)^k \\
&= \sum_k (z(1+z))^k \\
&= \frac{1}{1-z(1+z)} \\
&= \frac{1}{1-z-z^2}
\end{aligned}$$

This looks familiar! Indeed,

$$\begin{aligned}
f(n) &= < z^n > \frac{1}{1-z-z^n} \\
&= < z^{n+1} > \frac{z}{1-z-z^2} \\
&= F_{n+1}.
\end{aligned}$$

and – lo and behold! – the $(n+1)$st Fibonacci number!

We now turn to a somewhat harder problem. We want to crack:

$$\sum_k \binom{n+k}{m+2k} \cdot \binom{2 \cdot k}{k} \cdot \frac{(-1)^k}{k+1}$$

Notice we are here dealing with 2 parameters, n and m. We choose to treat m as fixed and set

$$f_m(n) := \sum_k \binom{n+k}{m+2k} \cdot \binom{2k}{k} \cdot \frac{(-1)^k}{k+1}.$$

Without a second thought, we form the generating function:

$$F_m(z) := \sum_n f_m(n) z^n.$$

Now, for the gymnastics , er, "generatingfunctionistics":

$$\begin{aligned}
F_m(z) &= \sum_n f_m(n) z^n \\
&= \sum_n \sum_k \binom{n+k}{m+2k} \cdot \binom{2k}{k} \cdot \frac{(-1)^k}{k+1} \cdot z^n \\
&= \sum_k \binom{2k}{k} \cdot \frac{(-1)^k}{k+1} \sum_n \binom{n+k}{m+2k} \cdot z^n \\
&= \sum_k \binom{2k}{k} \cdot \frac{(-1)^k}{k+1} \cdot z^{-k} \cdot \sum_n \binom{n+k}{m+2k} \cdot z^{n+k}
\end{aligned}$$

Here we're tempted to use the Binomial Theorem, but be careful! The exponent of z appears as the *upper* index of the binomial coefficient. For this, we offer as cure:

$$\begin{aligned}
\frac{x^m}{(1-x)^{m+1}} &= x^m (1-x)^{-(m+1)} \\
&= x^m \sum_k \binom{k+m}{k} x^k \\
&= \sum_k \binom{k+m}{m} x^{k+m} \\
&= \sum_k \binom{k}{m} x^k.
\end{aligned}$$

Now we can continue:

$$\begin{aligned}
F_m(z) &= \sum_k \binom{2k}{k} \cdot \frac{(-1)^k}{k+1} \cdot z^{-k} \cdot \frac{z^{m+2k}}{(1-z)^{m+2k+1}} \\
&= \frac{z^m}{(1-z)^{m-1}} \cdot \sum_k \binom{2k}{k} \frac{(-1)^k}{k+1} \cdot \frac{z^k}{(1-z)^{2k}} \\
&= \frac{z^m}{(1-z)^{m-1}} \cdot \sum_k \binom{2k}{k} \frac{1}{k+1} \cdot [\frac{-z}{(1-z)^2}]^k
\end{aligned}$$

Ah ha! We recognise the generating function for the Catalan numbers from last time:

$$\frac{1-\sqrt{1-4z}}{2z} = \sum_n C_n z^n = \sum_n \binom{2n}{n} \frac{1}{n+1} z^n.$$

So,

$$
\begin{aligned}
F_m(z) &= \frac{z^m}{(1-z)^{m-1}} \cdot \frac{1-\sqrt{1+\frac{4z}{(1-z)^2}}}{2\cdot(\frac{-z}{(1-z)^2})} \\
&= \frac{-z^{m-1}}{2\cdot(1-z)^{m-1}} \cdot [1-\sqrt{\frac{1+4z}{(1-z)^2}}] \\
&= \frac{-z^{m-1}}{2\cdot(1-z)^{m-1}} \cdot [1-\sqrt{\frac{(1+z)^2}{(1-z^2)}}] \\
&= \frac{-z^{m-1}}{2\cdot(1-z)^{m-1}} \cdot [1-\frac{1+z}{1-z}] \\
&= \frac{-z^{m+1}}{2\cdot(1-z)^{m-1}} \cdot \frac{(-2\cdot z)}{1-z} \\
&= \frac{z^m}{(1-z)^m}
\end{aligned}
$$

Hence, finally,

$$
\begin{aligned}
f_m(n) &= < z^n > F_m(z) \\
&= < z^n > \frac{z^m}{(1-z)^m} \\
&= < z^{n-1} > \frac{z^{m-1}}{(1-z)^m} \\
&= \binom{n-1}{m-1} \\
&= \binom{n-1}{n-m}.
\end{aligned}
$$

So finally, we have discovered that:

$$\sum_k \binom{n+k}{m+2k} \cdot \binom{2\cdot k}{k} \cdot \frac{(-1)^k}{k+1} = \binom{n-1}{n-m}.$$

Part of the philosophy of this course is to emphasise general methods and show how to apply them to particular problems. In this spirit, let us take another look at the problems we used as illustrations right on day 1. For this however, we will introduce yet more bizarre–looking things. Until now, we converted a sequence

$$a_0, a_1, a_2, \ldots$$

into the object

$$A(z) := \sum_n a_n z^n,$$

and then studied the latter. This is called the *ordinary* generating function of the sequence $a_0, a_1, a_2, \ldots$, in order to distinguish it from the *exponential generating function*

$$\hat{A}(z) := \sum_n a_n \frac{z^n}{n!}.$$

Let's see what happens when we multiply two of these together. Let

$$\hat{A}(z) := \sum_n a_n \cdot \frac{z^n}{n!}, \quad \hat{B}(z) := \sum_n b_n \cdot \frac{z^n}{n!},$$

be exponential generating functions for the sequences $a_0, a_1, a_2, \ldots$ and $b_0, b_1, b_2, \ldots$ respectively. Then

$$\begin{aligned}
\hat{A}(z) \cdot \hat{B}(z) &= (a_0 + a_1\frac{z}{1!} + a_2\frac{z^2}{2!} + a_3\frac{z^3}{3!} + \cdots)(b_0 + b_1\frac{z}{1!} + b_2\frac{z^2}{2!} + b_3\frac{z^3}{3!} + \cdots) \\
&= \frac{a_0}{0!}\frac{b_0}{0!} + (\frac{a_0}{0!}\frac{b_1}{1!} + \frac{a_1}{1!}\frac{b_0}{0!})z + (\frac{a_0}{0!}\frac{b_2}{2!} + \frac{a_1}{1!}\frac{b_1}{1!} + \frac{a_2}{2!}\frac{b_0}{0!})z^2 + \cdots
\end{aligned}$$

What is the coefficient of $\frac{z^n}{n!}$? It is

$$\begin{aligned}
\sum_{r+s=n} n!\frac{a_r}{r!}\frac{b_s}{s!} &= \sum_r \frac{n!}{r!(n-r)!} a_r b_{n-r} \\
&= \sum_r \binom{n}{r} a_r b_{n-r}.
\end{aligned}$$

This is the kind of convolution we get with exponential generating functions.

Let's apply this to the problem of derangements. There we were interested in the number of permutations $\pi : [n] \to [n]$ which have no *fixed points* i,e. the elements $i \in [n]$ such that $\pi(i) = i$. Let us classify the permutations of π according to their number of fixed points. A permutation that has k fixed points is obtained by picking a set of k elements, fixing them and taking a derangement on the remaining elements. From this we deduce the relation:

$$n! = \sum_k \binom{n}{k} d_{n-k}.$$

The right hand side is the exponential convolution of two sequences: the all 1 sequence and the sequence of derangements. The corresponding exponential generating functions are:

$$\sum_n 1 \cdot \frac{z^n}{n!} = e^z \quad \text{and} \quad \hat{D}(z) := \sum_n d_n \frac{z^n}{n!}.$$

The exponential generating function of the left hand side is

$$\sum_n n!\frac{z^n}{n!} = \sum_n z^n = \frac{1}{1-z}.$$

Hence we deduce the following equation relating the exponential generating functions:

$$\frac{1}{1-z} = e^z \hat{D}(z),$$

yielding

$$\hat{D}(z) = e^{-z}/(1-z).$$

Now we extract the coefficient of $z^n/n!$:

$$\begin{aligned} d_n &= <\frac{z^n}{n!}> \hat{D}(z) \\ &= n! < z^n > \hat{D}(z) \\ &= n! < z^n > \frac{e^{-z}}{1-z} \\ &= n! \sum_{k \leq n} < z^k > e^{-z} \\ &= n! \sum_{k \leq n} \frac{(-1)^k}{k!} \end{aligned}$$

which is the answer we got before!

Chapter 8

Probability Generating Functions

Recall the setting of a probability space $(\Omega, \Pr)$ where Ω is the sample space and the function $\Pr : \Omega \to [0,1]$ satisfies

$$\sum_{\omega} Pr(\omega) = 1.$$

An *event* is a subset $A \subseteq \Omega$, and by definition, has probability

$$\Pr(A) := \sum_{\omega \in A} \Pr(\omega).$$

Example 42 (Tossing a (fair) coin)

$$\Omega_1 = \{Heads, Tails\}$$

$$Pr_1(Heads) = \frac{1}{2} = Pr_1(Tails).$$

Example 43 (Tossing a (fair) dice)

$$\Omega_2 = \{side_1, side_2, side_3, side_4, side_5, side_6\}.$$

$$Pr_2(side_1) = \cdots = Pr_2(side_6) = \frac{1}{6}.$$

Interesting events are often specified via *random variables.* A random variable is a function $X : \Omega \to \mathsf{R}$. In the dice example one would choose $X(\text{side}_n) = n$. We could then define the event: "tossing at least a four",

$$A := \{\text{side}_4, \text{side}_5, \text{side}_6\} \equiv X \geq 4$$

$$Pr(A) = Pr(X \geq 4) = \frac{3}{6} = \frac{1}{2}.$$

One is often interested in summarising the behaviour of a random variable through certain *statistics*, the most basic of which is the *expectation* of X written $E[X]$, and sometimes as μ and defined by

$$\mu = E[X] = \sum_{\omega} X(\omega) \Pr(\omega). \tag{8.1}$$

In the situation where the random variable X is *discrete* i.e. only takes values $0, 1, 2, \ldots$, this can be equivalently written as:

$$\mu = E[X] = \sum_n n \cdot \Pr(X = n).$$

In the dice example this gives the following expectation,

$$E[X] = 1 \cdot \frac{1}{6} + 2 \cdot \frac{1}{6} + ... + 6 \cdot \frac{1}{6} = \frac{1}{6} \cdot \frac{6 \cdot 7}{2} = 3.5.$$

The expectation operator on random variables has the following nice *linearity* property: for random variables X, Y and reals α, β,

$$E[\alpha X + \beta Y] = \alpha E[X] + \beta E[Y].$$

The proof is a simple computation:

$$\begin{aligned}
E[\alpha X + \beta Y] &= \sum_\omega (\alpha X + \beta Y)(\omega) \Pr(\omega) \\
&= \sum_\omega (\alpha X(\omega) + \beta Y(\omega)) \Pr(\omega) \\
&= \sum_\omega \alpha X(\omega) \Pr(\omega) + \sum_\omega \beta Y(\omega) \Pr(\omega) \\
&= \alpha \sum_\omega X(\omega) \Pr(\omega) + \beta \sum_\omega Y(\omega) \Pr(\omega) \\
&= \alpha E[X] + \beta E[Y].
\end{aligned}$$

In the same way, we can extend this to:

$$E[\sum_i \alpha_i X_i] = \sum_i \alpha_i E[X_i].$$

In the following we will only be considering discrete random variables. In this situation, we are interested in the sequence:

$$p_k^X := \Pr(X = k),$$

and naturally we bring in the generating function

$$\begin{aligned}
G_X(z) &:= \sum_k p_k^X \cdot z^k && (8.2) \\
&= \sum_k \Pr(X = k) \cdot z^k. && (8.3)
\end{aligned}$$

This is called the *probability generating function* for the random variable X, and we shall abbreviate this to pgf. The pgf for the dice example is,

$$\begin{aligned}
G_2(z) &= \frac{1}{6} z + \frac{1}{6} z^2 + ... + \frac{1}{6} z^6 \\
&= \frac{1}{6}(z + z^2 + ... + z^6) \\
&= \frac{1}{6} \cdot \frac{1 - z^7}{1 - z}.
\end{aligned}$$

Note that the pgf of any random variable has the property that

$$G_X(1) = \sum_k \Pr(X = k) = 1. \tag{8.4}$$

If we know the pgf of a random variable, we can extract useful information about the statistics of the random variable. If we differentiate $G_X(z)$ we get,

$$G'_X(z) = \sum_k k \cdot \Pr(X = k) \cdot z^{k-1} \tag{8.5}$$

and hence,

$$G'_X(1) = \sum_k k \cdot \Pr(X = k) = E[X]. \tag{8.6}$$

The second most useful statistic of a random variable X, is the *variance*, written as $\sigma^2[X]$. It is defined by:

$$\sigma^2[X] := E[(X - E[X])^2].$$

Using the linearity property of expectations, we can derive an alterate equivalent expression as follows:

$$\begin{aligned}\sigma^2[X] := \quad & E[(X - E[X])^2] \\ = \quad & E[X^2 - 2E[X]X + (E[X])^2] \\ = \quad & E[X^2] - 2E[X]E[X] + E[(E[X])^2] \\ = \quad & E[X^2] - 2(E[X])^2 + (E[X])^2 \\ = \quad & E[X^2] - (E[X])^2.\end{aligned}$$

What is the significance of the variance? We shall see that it gives an idea about *how the random variable is distributed around its expectation.*

As with the expectation we would like to see if we can extract the variance out of the pgf of a random variable. From the second expression for the variance, it is clear that this entails extracting $E[X^2]$ out of the pgf. This parameter is called *the second moment* of X. By definition,

$$\begin{aligned}E[X^2] \quad &:= \quad \sum_\omega (X(\omega))^2 \Pr(\omega) \\ &= \quad \sum_k k^2 \cdot \Pr(X = k).\end{aligned}$$

The last equation holds, because X is a discrete random variable. To find the generating functions we start with

$$G(z) := \sum_k \Pr(X = k) z^k,$$

differentiate and multiply by z:

$$z \cdot G'(z) = \sum_k k \cdot \Pr(X = k) \cdot z^k.$$

If we differentiate this again we get,

$$G'(z) + z \cdot G''(z) = \sum_k k^2 \cdot \Pr(X = k) \cdot z^{k-1}.$$

Setting $z = 1$, we get:

$$\begin{aligned} G'(1) + G''(1) &= \sum_k k^2 \cdot \Pr(X = k) \\ &= E[X^2]. \end{aligned}$$

Thus we can extract the variance out of the pgf using the formula:

$$\sigma^2 = G'(1) + G''(1) - (G'(1))^2. \tag{8.7}$$

Given G_X, and G_Y, the pgfs of two random variables X and Y we would like to find the pgf of their sum $X + Y$. Let us compute:

$$\begin{aligned} G_{X+Y}(z) &:= \sum_k \Pr(X + Y = k) z^k \\ &= \sum_k \left(\sum_l \Pr(X = l, Y = k - l) \right) \cdot z^k \end{aligned}$$

To make further progress, we need an important assumption:

$$\Pr(X = x, Y = y) = \Pr(X = x) \cdot \Pr(Y = y),$$

for any two values $x, y \in reals$. In such a situation, the random variables X and Y are said to be *independent*. We expect this condition to hold when X and Y represent the outcomes of two different tosses of a coin, or two different throws of a dice. The independence assumption is a key one in probability theory, and for n random variables $X_1, \ldots, X_n$, it reads:

$$\Pr(X_1 = x_1, \cdots, X_n = x_n) = \Pr(X_1 = x_1) \cdot \cdots \cdot \Pr(X_n = x_n),$$

for any $x_1, \ldots, x_n \in \mathsf{R}$.

Under the assumption of independence, we can continue with our computation:

$$\begin{aligned} G_{X+Y}(z) &= \sum_k \left(\sum_l \Pr(X = l, Y = k - l) \right) \cdot z^k \\ &= \sum_k \left(\sum_l \Pr(X = l) \cdot \Pr(Y = k - l) \right) \cdot z^k \end{aligned}$$

Ah ha! The coefficient $\sum_l \Pr(X = l) \cdot \Pr(Y = k - l)$ rings a bell! Indeed, it is the convolution of the sequences of probabilities $\Pr(X = k)$ and $\Pr(Y = k)$. Thus remembering our rules for generating functions of convolutions,

$$\begin{aligned} G_{X+Y}(z) &= \sum_k \left(\sum_l \Pr(X = l) \cdot \Pr(Y = k - l) \right) \cdot z^k \\ &= \sum_k (< z^k > G_X(z) G_Y(z)) z^k \\ &= G_X(z) G_Y(z). \end{aligned}$$

Thus:

Theorem 44 *Let X and Y be two independent random variables with pgfs G_X and G_Y respectively. Then the pgf for their sum $X+Y$ is $G_{X+Y}(z) = G_X \cdot G_Y$. In general, if $X := X_1 + \cdots + X_n$, then $G_X(z) = G_{X_1}(z) \cdot \cdots \cdot G_{X_n}(z)$.*

Example 45 (Tossing a coin n times)

$$\Omega_3 = \{HH\cdots H, TH\cdots H, ..., TT\cdots T\} \quad , \quad |\Omega_3| = 2^n$$

One can think of the coin toss as modelling a random experiment. "Heads" represents "success" and "tails" "failure". Then we are interested in the random variable

$$X := \#successes = \#\{heads\ in\ the\ experiment\}.$$

We can decompose this into a sum of variables each relating to a single toss:

$$X = X_1 + X_2 + ... + X_n \quad , \quad where$$

$$X_i := \begin{cases} 1 & heads\ on\ the\ ith\ toss \\ 0 & else \end{cases}$$

So far we assumed a fair coin or dice, but let's be prepared for the casinos of Las Vegas and look at the general case, where,

$$\begin{aligned} \Pr(\text{Heads}) &= p \\ \Pr(\text{Tails}) &= q \\ p + q &= 1. \end{aligned}$$

This gives the following generating functions for the X_i's,

$$\begin{aligned} G_{X_i}(z) &= \sum_k \Pr(X_i = k) \cdot z^k \\ &= \Pr(X_i = 0) \cdot z^0 + \Pr(X_i = 1) \cdot z^1 \\ &= q + p \cdot z, \end{aligned} \tag{8.8}$$

and for X,

$$G_X(z) = (q + p \cdot z)^n. \tag{8.9}$$

Differentiating this gives,

$$G'_X(z) = p \cdot n \cdot (q + p \cdot z)^{n-1} \tag{8.10}$$

$$G''_X(z) = p^2 \cdot n \cdot (n-1) \cdot (q + p \cdot z)^{n-2}. \tag{8.11}$$

We are now able to compute the expectation and the variance for the n coin example,

$$E[X] = G'_X(1) = p \cdot n \tag{8.12}$$

$$\begin{aligned} \sigma^2[X] &= G'_X(1) + G''_X(1) - (G'_X(1))^2 \\ &= p \cdot n + p^2 \cdot n \cdot (n-1) - (p \cdot n)^2 \\ &= p \cdot n \cdot (1 - p) \\ &= npq. \end{aligned} \tag{8.13}$$

We hinted earlier that the variance of a random variable gives an indication of how the random variable is distributed around the expectation. The following famous theorem makes this precise.

Theorem 46 (Chebyshev's Inequality) *Let X be a random variable with expectation μ and variance σ^2. Then for all $t > 0$,*

$$\Pr(|X - \mu| > t \cdot \sigma) \leq \frac{1}{t^2} \tag{8.14}$$

Proof:

$$\begin{aligned}
\sigma^2 &:= E[(X-\mu)^2] \\
&= \sum_{\omega} (X(\omega) - \mu)^2 \cdot \Pr(\omega) \\
&\geq \sum_{\omega:|X(\omega)-\mu|>t\cdot\sigma} (X(\omega)-\mu)^2 \cdot \Pr(\omega) \\
&\geq \sum_{\omega:|X(\omega)-\mu|>t\cdot\sigma} (t\cdot\sigma)^2 \cdot \Pr(\omega) \\
&= t^2 \cdot \sigma^2 \cdot \sum_{\omega:|X(\omega)-\mu|>t\cdot\sigma} \Pr(\omega) \\
&= t^2 \cdot \sigma^2 \cdot \Pr(|X-\mu| > t\cdot\sigma).
\end{aligned}$$

Dividing by σ^2 completes the proof. □

Let's go back to the n coin example. Assume it's a fair coin situation, and let n be 1000.

$$E[X] = \frac{1}{2} \cdot 1000 = 500$$

$$\sigma^2[X] = \frac{1}{2} \cdot 1000 \cdot \frac{1}{2} = 250$$

$$\Pr(|X - 500| > t \cdot \sqrt{250}) < \frac{1}{t^2}$$

$$\Pr(|X - 500| > t \cdot 20) < \frac{1}{t^2}$$

Let $t = 2$,

$$\Pr(|X - 500| > 40) < \frac{1}{4}$$

$$\Pr(460 \leq X \leq 540) \geq \frac{3}{4}.$$

Example 47 *Toss a coin with $\Pr(\mathit{Heads}) = p$ and $\Pr(\mathit{Tails}) = q$ until you get "heads". How many tosses do we need? The sample space is*

$$\Omega_4 = \{H, TH, TTH, TTTH, ...\}$$

which is different from our previous examples because it is infinite! The random variable of interest is $X := \#\{\text{tosses before you get a head}\}$ and the probability measure $\Pr_4$ is prescribed by assuming that the tosses are all independent of each other: $\Pr_4(X = k) = q^{k-1}p$. for $k \geq 1$. (Of course, $\Pr(X = 0) = 0$!) The pgf for this random variable is:

$$G_X(z) = \sum_k \Pr(X = k) \cdot z^k$$

$$
\begin{aligned}
&= \sum_{k\geq 1} q^{k-1}\cdot p\cdot z^k \\
&= pz\sum_{k\geq 0} q^k z^k \\
&= \frac{pz}{1-qz}.
\end{aligned}
\tag{8.15}
$$

Check that $G(1) = p/(1-q) = 1$ as it should. What is the expected number of tosses befor the first "heads"? We have

$$G'_X(z) = \frac{p}{(1-qz)^2},$$

hence

$$E[X] = G'_X(1) = 1/p.$$

Thus for a fair coin, $p = 1/2$, we expect to wait for two tosses before getting a "heads".

Exercise 48 *By computing the variance, give more precise information on the distribution of this random variable. make the computation also for the random variable denoting the number of trials before we have k "Heads" for $k \geq 1$.*

Chapter 9

A coin flipping game

In this section we shall use generating functions to analyse Penney's coin flipping game. The game is played by two players A and B. A starts by choosing a sequence of l heads and tails. Then, B chooses a different sequence of l head and tails, possibly relying on A's choice. A fair coin is flipped until either A's or B's sequence comes up in the sequence of coin tosses. The winner is the player whose sequence appears first.

We begin by analysing the game when the two players, the class A and the lecturer B, choose sequences of length one, two or three. We give the class A the privilege of choosing first and the lecturer B, will choose the second sequence.

Length 1 No matter whether A chooses H or T, B must choose the opposite. The game is over after one coin flip and both players win with probability $\frac{1}{2}$.

Length 2 Suppose A chooses HH; then B can choose TH. A wins only if the first two coin tosses come out H. That is, A wins with probability $\frac{1}{4}$ and B with probability $\frac{3}{4}$. If A instead chooses TH, B can choose either HT or TT, both choices with equal chances for A and B.

Length 3 Suppose A chooses HTH. Then B may choose HHT. A and B both need an H at the beginning of their sequences so they start by flipping the coin until an H comes up. We can split the analysis of the following two coin tosses into the following cases:

HH: B wins, since a T eventually will show up.

HT: B wins.

TH: A wins.

TT: The players start over again waiting for an H.

So B wins two times out of three giving A chances $\frac{1}{3}$ against B's $\frac{2}{3}$ probability of winning.

If A tried to pick B's last choice HHT, B can choose THH. In this case A wins if and only if the first two flips both come up heads giving A probability $\frac{1}{4}$ of winning against B's $\frac{3}{4}$.

If A tried to choose THH, B could choose TTH, which by similar arguments as used above again gives B probability $\frac{2}{3}$ of winning.

By enumerating all 8 sequences of three head and tails we can check that for any choice of A, B can choose a sequence that has strictly higher probability of winning, or in other words, a "better" sequence than A's. Penney's coin flipping game is a *non-transitive* game in the sense that the "better" relation is not transitive.

Definition 49 *Let A and B be two sequences over an alphabet Ω, and let $n = |A|$ and $q = |\Omega|$. AB denotes the* correlation *between A and B which is defined to be a sequence in $\{0,1\}^n$, illustrated by the following example. Let $A =$ `abcabca` and $B =$ `abcabcd`.*

$$
\begin{array}{lllllllllllllll}
A & = & a & b & c & \mathbf{a} & \mathbf{b} & \mathbf{c} & \mathbf{a} & & & & & & AB \\
B & = & a & b & c & a & b & c & d & & & & & & 0 \\
 & & & a & b & c & a & b & c & d & & & & & 0 \\
 & & & & a & b & c & a & b & c & d & & & & 0 \\
 & & & & & \mathbf{a} & \mathbf{b} & \mathbf{c} & \mathbf{a} & b & c & d & & & 1 \\
 & & & & & & a & b & c & a & b & c & d & & 0 \\
 & & & & & & & a & b & c & a & b & c & d & 0 \\
 & & & & & & & & \mathbf{a} & b & c & a & b & c\ \ d & 1
\end{array}
$$

The correlation is $AB = 0001001$ as shown to the right and is calculated in the following way. First place B below A. Then shift B one position to right $n-1$ times. At each shifted B consider the symbols exactly below A's symbols. If all these symbols are equal, then AB has 1 at the corresponding index and otherwise 0. In the above example bold symbols indicates this.

Formally for $A = A^0 \cdots A^{n-1}$ and $B = B^0 \cdots B^{m-1}$, the correlation $AB = AB^0 \cdots AB^{n-1}$ is defined by:

$$AB^i = \begin{cases} 1 & \text{if } A^{i+j} = B^j, \text{ for all } 0 \le j < \min(n-i, m) \\ 0 & \text{otherwise.} \end{cases}$$

Observe that the correlation is not symmetric. Namely, AB is not necessarily equal to BA. The correlation AA is called A's autocorrelation.

We are also going to express the correlation AB in another way: AB_z denotes the integer AB viewed as a number in base z, with the least significant digit on the right (AB^{n-1}). Alternatively, AB_z can be viewed as the polynomial in the variable z with coefficients AB^i. Namely:

$$AB_z = \sum_{i=1}^{n} AB^{n-i} z^{i-1}$$

Let's generalize the game above so that instead of flipping a coin we use a fair dice with q sides. Denote by Ω the set of $|\Omega| = q$ possible outcomes of throwing the dice. Now A and B are chosen to be sequences over the alphabet Ω. The following theorem gives the formula for the *odds* in Penney's game.

Theorem 50 (Conway) *The* odds *in favour of B are:*

$$\frac{Pr(B\ wins)}{Pr(A\ wins)} = \frac{AA_q - AB_q}{BB_q - BA_q}.$$

To prove this theorem it will be necessary to introduce some machinery that we shall develop using generating functions.

Let A be a string over the alphabet Ω. The functions $f, f_A : \mathsf{N}_0 \to \mathsf{N}_0$ are defined relative to A as follows.

$$
\begin{aligned}
f(n) &= \text{number of sequences in } \Omega^n \text{ not containing } A. \\
f_A(n) &= \text{number of sequences in } \Omega^n \text{ that have a } \textit{single} \text{ occurrence of } A, \text{ at their end.}
\end{aligned}
$$

The corresponding generating functions F and F_A are defined as:

$$
\begin{aligned}
F(z) &= \sum_{n=0}^{\infty} f(n) z^{-n} \\
F_A(z) &= \sum_{n=0}^{\infty} f_A(n) z^{-n}
\end{aligned}
$$

Notice that these generating functions are different than the ones we have seen in the previous sections as the exponents of z are negative. However, these generating functions can be treated exactly in the same way as those with positive exponents. The negative exponents simplify our notation. You are encouraged to work out the whole analysis below, changing the definition of these generating functions to use positive exponents and see the difference!

As we have seen before, generating functions are especially useful when we are able to express some recursive relation on the coefficients. We shall find next such relations for f and f_A. We first observe that $f_A(n) = 0$, for $n < |A|$, and $f(0) = 1$.

Consider a string of length n without occurrences of A. If we extend it with one symbol at the end there are two possible outcomes: either the extended string ends with A or not. In the latter case it is counted as one of the $f(n+1)$ string not containing A and in the former case it is one of the $f_A(n+1)$ strings ending with A, but containing no other occurrence of A. Since we can extend the string with q different symbols we have:

$$
q\, f(n) = f(n+1) + f_A(n+1).
$$

Multiplying by z^{-n} and summing for $n = 0, \ldots, \infty$, we get:

$$
q \sum_{n=0}^{\infty} f(n) z^{-n} = \sum_{n=0}^{\infty} f(n+1) z^{-n} + \sum_{n=0}^{\infty} f_A(n+1) z^{-n} \tag{9.1}
$$

Using the fact that $f(0) = 1$ we have

$$
\begin{aligned}
\sum_{n=0}^{\infty} f(n+1) z^{-n} &= z \sum_{n=0}^{\infty} f(n+1) z^{-(n+1)} \\
&= z \sum_{n=1}^{\infty} f(n) z^{-n} \\
&= z(F(z) - f(0)) = zF(z) - z
\end{aligned} \tag{9.2}
$$

and since $f_A(0) = 0$ we similarly have

$$
\sum_{n=0}^{\infty} f_A(n+1) z^{-n} = z(F_A(z) - f_A(0)) = zF_A(z) \tag{9.3}
$$

By (9.1), (9.2) and (9.3) we get the generating function equation:

$$qF(z) = zF(z) - z + zF_A(z)$$
$$\Downarrow$$
$$(z-q)F(z) + zF_A(z) = z \tag{9.4}$$

We can relate f and f_A yet in another way:

$$f(n) = \sum_{r \in AA} f_A(n+r), \tag{9.5}$$

where here, the autocorrelation AA is the set of indices $\{i \mid AA^{|A|-i} = 1\}$. (We count the indices from the right starting with 1.) We also have the identity

$$\sum_{i \in AA} z^i = zAA_z. \tag{9.6}$$

We now prove equation (9.5). Fix $n \geq 0$ and let M be the set of strings in Ω^n not containing A. That is $|M| = f(n)$. Similarly let M_A be the set of strings in $\bigcup_{r \in AA} \Omega^{n+r}$ with a *single* occurrence of A at the end. Then, $|M_A| = \sum_{r \in AA} f_A(n)$. Consider the mapping $\pi : M \to M_A$ defined to be:

$$\pi(W) = \text{prefix of } WA \text{ of least length with } A \text{ at its end.}$$

The function $\pi(W)$ is well-defined since any such WA contains A and its prefix of least length ending with A is unique. The prefix $\pi(W)$ must have length $n+r$, for some $r \in AA$, since the prefix has to be longer than W and it must end with an occurrence of A overlapping the last occurrence of A in WA.

We claim that π is bijective. It clearly is 1-1 since $\pi(W) = \pi(W') \Rightarrow W = W'$. Suppose that π is not onto. Let $W_A \in M_A$ be a string such that $\forall S \in M : \pi(S) \neq W_A$ and let $r \in AA$ be such that $|W_A| = n + r$. Let W be the prefix of W_A of length n. Consider a prefix of WA of length $n+r$. Since $r \in AA$, this prefix ends with A, and must thus equal W_A. But since W_A only has one occurrence of A it must be the shortest prefix of WA ending with A. Hence, $\pi(W) = W_A$; a contradiction.

Since π is a bijection between the finite sets M and M_A, they have the same cardinality, from which (9.5) follows. Multiplying by z^{-n} an summing over $n = 0, \ldots, \infty$, we get that

$$F(z) = \sum_{n=0}^{\infty} f(n) z^{-n} = \sum_{n=0}^{\infty} \Big(\sum_{r \in AA} f_A(n+r) \Big) z^{-n}.$$

Since $1 \leq r \leq |A|$, and $f_A(i) = 0$, for $0 \leq i < |A|$, we obtain similarly to (9.3) that:

$$\sum_{n=0}^{\infty} f_A(n+r) z^{-n} = z^r F_A(z)$$

and hence, by (9.6),

$$F(z) \;=\; \sum_{n=0}^{\infty} \Big(\sum_{r \in AA} f_A(n+r) \Big) z^{-n}$$

$$
\begin{aligned}
&= \sum_{r\in AA}\sum_{n=0}^{\infty} f_A(n+r)z^{-n} \\
&= \sum_{r\in AA} z^r F_A(z) \\
&= zAA_z F_A(z). \qquad (9.7)
\end{aligned}
$$

Solving (9.4) and (9.7) simultaneously, we obtain the following formulae:

$$
F_A(z) = \frac{1}{1+(z-q)AA_z} \qquad (9.8)
$$

$$
F(z) = \frac{zAA_z}{1+(z-q)AA_z} \qquad (9.9)
$$

Lets calculate the expected waiting time for a sequence $A \in \Omega^n$ to come up. Let X be a random variable, such that $X = k$ if we see an occurrence of the sequence A for the first time after k coin tosses and not before. We want to calculate the expected waiting time for A, which is the expectation $E[X] = \sum_{i=0}^{\infty} Pr(X = i)i$.

We shall consider two ways to calculate $E[X]$. In the first we shall use $f(n)$. Observe that the probability:

$$
Pr(X > n) = \frac{f(n)}{q^n},
$$

and

$$
Pr(X = n) = Pr(X > n-1) - Pr(X > n).
$$

We then get that:

$$
\begin{aligned}
E[X] &= \sum_{n=0}^{\infty} Pr(X = n)n \\
&= \sum_{n=1}^{\infty} (Pr(X > n-1) - Pr(X > n))n \\
&= Pr(X > 0) \quad - Pr(X > 1) + 2Pr(X > 1) \quad - 2Pr(X > 2) + 3Pr(X > 2) - \cdots \\
&= \sum_{n=0}^{\infty} Pr(X > n) \\
&= \sum_{n=0}^{\infty} \frac{f(n)}{q^n} \\
&= F(q) \;=\; qAA_q
\end{aligned}
$$

We can also use $f_A(n)$ to calculate this expectation, observing that:

$$
Pr(X = n) = \frac{f_A(n)}{q^n}.
$$

Recall that

$$
F_A(z) = \sum_{n=0}^{\infty} f_A(n)z^{-n}.
$$

Taking the derivative we get:

$$F_A'(z) = -\sum_{n=0}^{\infty} n f_A(n) z^{-n-1}.$$

Multiplying by $-z$, we get

$$-zF_A'(z) = \sum_{n=0}^{\infty} n f_A(n) z^{-n}.$$

Which is exactly what we want if we substitute z by q. Now,

$$F_A'(z) = -\frac{AA_z + (z-q)AA_z'}{(1+(z-q)AA_z)^2}$$

and

$$E[X] = (-zF_A'(z))_{|z=q} = qAA_q.$$

Exercise 51 *What is the variance of X?*

Consider the sequences HTT and THH. In the examples at the beginning of the section we have seen that THH has higher probability of winning against HTT. However, the auto-correlation of both strings are identical, and hence, the expected waiting times for both strings are identical. This shows that by comparing the waiting times we do not get an indication of which sequence in "better" in Penney's game.

Now let's generalize this analysis to more than one sequence. Let $\{A, B, C, \ldots, T\}$ be a *reduced* set of strings over the alphabet Ω, or in other words, a set in which no string is a substring of another. The function $f : \mathsf{N}_0 \to \mathsf{N}_0$ is defined relative to $A, B, \ldots, T$, so that:

$$f(n) = \text{number of sequences in } \Omega^n \text{ not containing any } A, B, \ldots, T.$$

The functions $f_H : \mathsf{N}_0 \to \mathsf{N}_0$ are defined for all $H \in \{A, \ldots, T\}$, relative to $A, B, \ldots, T$, so that:

$$f_H(n) = \begin{array}{l}\text{number of sequences in } \Omega^n \text{ not containing } A, \ldots, T, \\ \text{except a } \textit{single} \text{ occurrence of } H, \text{ at their end.}\end{array}$$

The corresponding generating functions F and F_H, for $H \in \{A, \ldots, T\}$ are defined as:

$$\begin{aligned} F(z) &= \sum_{n=0}^{\infty} f(n) z^{-n} \\ F_H(z) &= \sum_{n=0}^{\infty} f_H(n) z^{-n} \end{aligned}$$

We now establish some relations between these functions, similarly to the relations established above for one string. Given a string of length n without any occurrence of $A, \ldots, T$, it can be extended to a string of length $n+1$ by appending one of the q different alphabet symbols at its end. The resulting string might contain no occurrence of $A, \ldots, T$, or a single occurrence of one of the strings $A, \ldots, T$ at its end. This gives us the following equality:

$$q\, f(n) = f(n+1) + f_A(n+1) + f_B(n+1) + \ldots + f_T(n+1).$$

Multiplying by z^{-n} and summing for $n = 0, \ldots, \infty$, we get:

$$q \sum_{n=0}^{\infty} f(n) z^{-n} = \sum_{n=0}^{\infty} f(n+1) z^{-n} + \sum_{n=0}^{\infty} f_A(n+1) z^{-n} + \cdots + \sum_{n=0}^{\infty} f_T(n+1) z^{-n}. \quad (9.10)$$

By (9.10), (9.2) and (9.3) we have

$$\begin{gathered} qF(z) = zF(z) - z + zF_A(z) + \cdots + zF_T(z) \\ \Downarrow \\ (z-q)F(z) + zF_A(z) + \cdots + zF_T(z) = z \end{gathered} \quad (9.11)$$

Similarly to the analysis above for a single string (9.5), for any $H \in \{A, \ldots, T\}$, we establish the relation:

$$f(n) = \sum_{r \in AH} f_A(n+r) + \sum_{r \in BH} f_B(n+r) + \cdots + \sum_{r \in TH} f(n+r).$$

Multiplying by z^{-n} an summing over $n = 0, \ldots, \infty$, we get that

$$\sum_{n=0}^{\infty} f(n) z^{-n} = \sum_{n=0}^{\infty} \Big(\sum_{r \in AH} f_A(n+r) \Big) z^{-n} + \cdots + \sum_{n=0}^{\infty} \Big(\sum_{r \in TH} f_A(n+r) \Big) z^{-n}.$$

Which is the generating function equality:

$$F(z) = zAH_z F_A(z) + \cdots + zTH_z F_T(z). \quad (9.12)$$

Equations (9.11) and (9.12), for $H \in \{A, \ldots, T\}$ can be summarized as:

$$\begin{pmatrix} z-q & z & \cdots & z \\ 1 & -zAA_z & \cdots & -zTA_z \\ & \cdots & \cdots & \\ 1 & -zAT_z & \cdots & -zTT_z \end{pmatrix} \begin{pmatrix} F(z) \\ F_A(z) \\ \cdots \\ F_T(z) \end{pmatrix} = \begin{pmatrix} z \\ 0 \\ \cdots \\ 0 \end{pmatrix}.$$

Since the set $\{A, \ldots, T\}$ is a reduced set of strings, in each column of the matrix above the largest degree polynomial occurs on the diagonal and the only other polynomial that can have equal degree is in the first row. Hence, by the expansion of the determinant of the matrix, the determinant is a non-zero polynomial (of degree $1 + |A| + |B| + \cdots + |T|$) and the solution of the system of equations above is unique.

We now use this formulae in Penney's game.

Proof:(of Theorem 50) Let A and B be two different strings of equal length. Then the set $\{A, B\}$ is a reduced set of strings. By the analysis above we get the system of equations:

$$\begin{aligned} (z-q)F(z) + zF_A(z) + zF_B(z) &= z \\ F(z) + zAA_z F_A(z) + zBA_z F_B(z) &= 0 \\ F(z) + zAB_z F_A(z) + zBB_z F_B(z) &= 0. \end{aligned}$$

Solving the system we get that:

$$F_A(z) = \frac{BB_z - BA_z}{(z-q)(AA_zBB_z - AB_zBA_z) + AA_z + BB_z - AB_z - BA_z}$$

and by symmetry the expression for $F_B(z)$ is the same – exchanging A and B. Now we can finally determine the probability that A wins:

$$P(A \text{ wins}) = \sum_{n=0}^{\infty} f_A(n)q^{-n} = F_A(q) = \frac{BB_q - BA_q}{AA_q + BB_q - AB_q - BA_q}.$$

The probability the B wins is of course symmetric. We now get Conway's fomula for the odds in favour of B:

$$\frac{P(B \text{ wins})}{P(A \text{ wins})} = \frac{AA_q - AB_q}{BB_q - BA_q},$$

completing the proof. $\square$

We next prove that Penney's game is non-transitive:

Theorem 52 *Given any string A over an alphabet Ω, such that $|A| \geq 3$ and $|\Omega| \geq 2$ or $|A| \geq 2$ and $|\Omega| \geq 3$, then there exists a string B of the same length with the odds on B's side.*

Proof: Given $A = a_1a_2\ldots,a_l$ let $B = ba_1a_2\ldots a_{l-1}$, with $b \neq a_2$. We claim then that the odds

$$\frac{P(B \text{ wins})}{P(A \text{ wins})} = \frac{AA_q - AB_q}{BB_q - BA_q} > 1.$$

We shall only prove this here for $|\Omega| = 2$ and $|A| \geq 3$ and leave the general case, that can be proved similarly, as an exercise.

We claim that $AA_2 - AB_2 \geq 2^{l-2} + 2^{l-3}$ and that $BB_2 - BA_2 < 2^{l-2} + 2^{l-3}$. To see this, consider the binary expansion of AA_2 and AB_2. The most significant digit of AA_2 is 1, while the two most significant digits of AB_2 are 00. In addition, apart from the least significant digits of AA_2 and AB_2, it is not possible that both AA_2 and AB_2 have a 1 in the same position since this implies that $a_1 = b = a_2$. Also, if some digit, other than the least significant, is 0 in AA_2 and 1 in AB_2, then the digit on its right is 1 in AA_2 and 0 in AB_2. All this implies that $AA_2 - AB_2 \geq 2^{l-2} + 2^{l-3}$, since if the difference borrows from the most significant digit of AA_2, then it leaves a 1 in the next two most significant digits.

Consider now the binary expansion of BB_2 and AB_2. The three most significant digits of BB_2 are 100 since $b \neq a_2$. The two most significant digits of BA_2 are 01. Hence, $BB_2 - BA_2 < 2^{l-2} + 2^{l-3}$, since if the difference borrows from the second most significant digit of BB_2, then the difference is smaller than 2^{l-2} and if it does not borrow, then it is smaller than $2^{l-2} + 2^{l-3}$. $\square$

Part II

Graph Theory

Chapter 10

Basics

The following lectures will be on Graph Theory.

We assume that the students are familiar with the basic concepts, definitions, and some theorems that are given in mandatory courses at Aarhus. For example, the Max-Flow-Min-Cut Theorem.

Section 19 contains some problems and their solutions.

Since we are all familiar with the basic concepts and the definitions, we shall start with a quick reminder and introduce the notation that we will use throughout this chapter.

10.1 Graphs

A *graph* G consists of a pair (V, E) of sets of *vertices* V and *edges* E. An edge $e \in E$ is a pair $e = (v_1, v_2)$ *connecting* the two vertices $v_1, v_2 \in V$. A graph can be *directed* (called also *digraph*) when its edges have a "direction" or in other words are ordered pairs; otherwise the edges are unordered pairs and the graph is *undirected*. Two vertices can be connected with multiple edges, in which case E is a multi-set. A graph with multiple edges is called a *multiple-edged* graph or *multigraph*. A *loop* is an edge connecting a vertex to itself, e.g. $(v, v) \in E$. A *simple* graph is a graph without loops and without multiple edges. In these lectures we will mainly be studying properties of undirected graphs.

Two vertices v_1 and v_2 are *adjacent* in G if they are connected by an edge, e.g. $(v_1, v_2) \in E$. An edge $(v_1, v_2) \in E$ is *incident* to the vertices v_1 and v_2. Given a graph $G = (V, E)$ and an edge e we denote by $G + e$ the graph $(V, E \cup \{e\})$ by $G - e$ the graph $(V, E \setminus \{e\})$, etc. The sets of edges and vertices of a graph G will sometimes be denoted by $E(G)$ and $V(G)$.

10.2 Subgraphs

A *subgraph* G' of G is a graph $G' = (V', E')$ with $V' \subseteq V$ and $E' \subseteq E$. G' is a *spanning* subgraph if $V' = V$. G' is an *induced* subgraph if E' contains every edge in E that has both its endpoints in V'. We will use the notation $G[V']$ to denote the subgraph of G that is induces by the set of vertices $V' \subseteq V$.

10.3 Isomorphism

Two simple undirected graphs $G = (V_1, E_1)$ and $H = (V_2, E_2)$ are *isomorphic*, denoted $G \cong H$, if and only if there exist a bijective mapping $f : V_1 \to V_2$ such that

$$(v_1, v_2) \in E_1 \Leftrightarrow (f(v_1), f(v_2)) \in E_2.$$

An *automorphism* of $G = (V, E)$ is a permutation σ of V such that

$$(v_1, v_2) \in E \Leftrightarrow (\sigma(v_1), \sigma(v_2)) \in E.$$

For multigraphs these definitions can be extended such that the number of edges (v_1, v_2) in the multi-set E_1 has to be equal to the number of edges $(f(v_1), f(v_2))$ in E_2. Similar definitions can be given for directed graphs. See Figure 10.1.

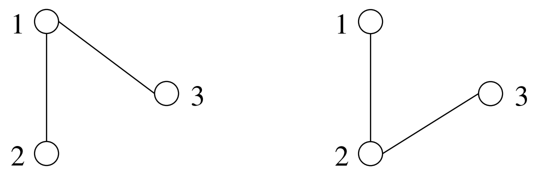

Figure 10.1: Two isomorphic graphs . Each has two automorphisms: the identity permutation on the vertices and the permutation that swaps the two non-adjacent vertices.

Exercise 53 *Prove that there are exactly* $2n$ *automorphisms of a cycle with* n *vertices,* $n \geq 3$.

10.4 Incidence and adjacency matrices

A graph $G = (V, E)$ with n vertices $V = \{v_1, \ldots, v_n\}$ and m edges $E = \{e_1, \ldots, e_m\}$ can be represented by two types of matrices; the *adjacency matrix* and the *incidence matrix.*

The adjacency matrix is an $n \times n$ matrix M. For an undirected graph the (i, j)'th entry $M(i, j)$ equals the number of edges connecting v_i with v_j; for directed graphs the entry is the number of edges directed from v_i to v_j. Loops at v_i contribute 2 to $M(i, i)$.

$$\begin{array}{c c}
 & \begin{array}{cccc} v_1 & v_2 & \ldots & v_n \end{array} \\
\begin{array}{c} v_1 \\ v_2 \\ \vdots \\ v_n \end{array} &
\left[\begin{array}{cccc} 0 & 1 & \ldots & 0 \\ 2 & 1 & \ldots & 1 \\ \vdots & \vdots & & \vdots \\ 0 & 1 & \ldots & 0 \end{array}\right]
\end{array}$$

Figure 10.2: Example of an adjacency matrix.

The incidence matrix is a $n \times m$ matrix with a row for each vertex and a column for each edge. $M(i, j)$ is 1 if there is an edge e_j incident to v_i and zero otherwise, with the exception that if e_j is a loop then $M(i, j) = 2$.

$$\begin{array}{c} \\ v_1 \\ v_2 \\ \vdots \\ v_n \end{array} \begin{array}{c} \begin{array}{cccc} e_1 & e_2 & \dots & e_m \end{array} \\ \begin{bmatrix} 0 & 1 & \dots & 0 \\ 2 & 0 & \dots & 1 \\ \vdots & \vdots & & \vdots \\ 0 & 1 & \dots & 1 \end{bmatrix} \end{array}$$

Figure 10.3: Example of an incidence matrix.

Fact 54 *There are $2^{\binom{n}{2}}$ simple labeled undirected graphs. (By labeled we mean that we distinguish between isomorphic graphs that are not equal.)*

10.5 Paths and cycles

A *walk* in a graph $G = (V, E)$ is a sequence $v_1 e_1 v_2 e_2 \cdots e_{k-1} v_k$ of vertices $v_i \in V$ and edges $e_i \in E$. The edges e_i are included in case of a non-simple graph with multiple edges. The walk is of *length* k. A walk with distinct *edges* is called a *trail*. Similarly a *path* is a walk with distinct *vertices*. A *closed walk* is a walk starting and ending at the same vertex, e.g. $v_1 = v_k$. Similarly a *circuit* is a "closed trail" and a *cycle* a "closed path" with all vertices except the first and last being different. A graph is *acyclic* if it does not have any cycles. Two vertices u, v are *connected* if there exists a path from u to v and a graph is *connected* if all vertices are pairwise connected.

Analogous definitions for digraphs require that the edges on a walk (trial, path, etc.) will all be directed in the same direction. Now, a vertex x is connected to a vertex y if there is a directed path from x to y. A digraph is *strongly connected* if for every pair of vertices x and y, x is connected to y and y is connected to x.

10.6 Degree

For an undirected graph the *degree* of a vertex v is the number of edges incident to v, counting loops twice. We denote the degree of v by $d(v)$. For directed graphs there are an *in-degree* and an *out-degree*. The in-degree $d^-(v)$ of a vertex v is the number of edges of the form $\{(v', v) \in E \mid v' \in V\}$, and similarly, the out-degree $d^+(v)$ of v is the number of edges in $\{(v, v') \in E \mid v' \in V\}$.

Theorem 55 *Let $G = (V, E)$ be an undirected graph. Then*

$$\sum_{v \in V} d(v) = 2 \cdot |E|.$$

Proof: Consider the incidence matrix M of G. The sum of the entries in each column of M is 2. On the other hand, the sum of the entries in each row equals the degree of the corresponding vertex, and thus we have

$$\sum_{v \in V} d(v) = \sum_{i=1}^{|V|} \sum_{j=1}^{|E|} M(i,j) = \sum_{j=1}^{|E|} \sum_{i=1}^{|V|} M(i,j) = \sum_{j=1}^{|E|} 2 = 2 \cdot |E|. \quad \square$$

Corollary 56 *An undirected graph G has an even number of vertices with odd degree.*

Proof:

$$\sum_{v\in V} d(v) = \sum_{v:d(v)\text{ is odd}} d(v) + \sum_{v:d(v)\text{ is even}} d(v).$$

By Theorem 55, the left-hand-side of the equation above must be an even integer. Thus, the first term of the right-hand-side must be even implying the corollary. □

Arguments like the one used in the last proof are called *parity arguments*.

Given an undirected graph $G = (V, E)$, define the minimum degree $\delta(G) = \min\{d(v) \mid v \in V\}$ and the maximum degree $\Delta(G) = \max\{d(v) \mid v \in V\}$. G is called *k-regular* if $d(v) = k$ for all $v \in V$, or in other words if $\delta(G) = \Delta(G)$. Figures 10.4, 10.5 and 10.6 give examples of regular graphs.

Exercise 57 *There exist a k-regular simple graph with n vertices if and only if either k or n are even (or both) and $n \geq k + 1$.*

Solution: Let k and n be fixed. Clearly, if the graph is to be simple, then $n \geq k + 1$. If both k and n are odd, then a k-regular graph with n vertices will contradict Corollary 56. Therefore, either n or k must be even. Suppose that k is even. We can line up n vertices on a circle and counting clockwise we connect each vertex with the $\frac{k}{2}$ subsequent vertices. Now since $\frac{k}{2} < \frac{n}{2}$ every vertex v is connected to $\frac{k}{2}$ distinct vertices before and after itself on the circle, each contributing one edge incident to v implying that v has degree $\frac{k}{2} + \frac{k}{2} = k$. When k is odd (and n even) start with the construction above for $k - 1$. For a vertex v the $\frac{n}{2}$*'th* vertex starting from v (the opposite vertex on the circle) has not been connected to v. So by connecting all vertices with their opposite vertices the claim is satisfied. See Figure 10.4.

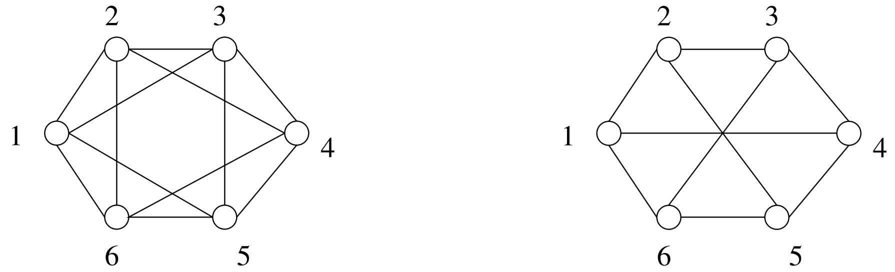

Figure 10.4: Examples of a 4-regular and 3-regular graphs with 6 vertices.

10.7 Special graphs

We will refer to some special graphs in these lectures. The *complement* G^c of a simple graph $G = (V, E)$ is the graph $G^c = (V, E')$ where $(u, v) \in E'$ if and only if $(u, v) \notin E$. The *complete* graph, also called a *clique*, is a graph in which every two vertices are adjacent. We denote a complete graph with n vertices by K_n. An *independent* set of vertices is the complement of a clique in the sense that none of the vertices in the independent set are connected by an edge. A graph $G = (V, E)$ is *bipartite* if there is a partition of the vertices V into two disjoint independent subsets X and Y, $V = X \cup Y$. A *complete bipartite* graph $G = (V, E)$ is a bipartite graph such that for any two vertices $u \in X$ and $v \in Y$, the edge $(u, v) \in E$, when X and Y are the bi-partition of V. The complete bipartite graph with $|X| = n$ and $|Y| = m$ vertices is denoted by $K_{n,m}$. See Figures 10.5 and 10.6.

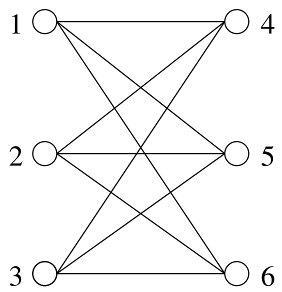

Figure 10.5: The complete bipartite graph $K_{3,3}$.

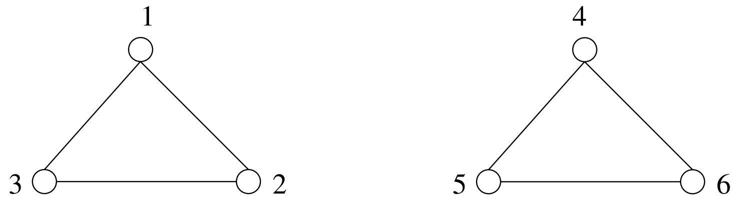

Figure 10.6: The complement of $K_{3,3}$ is formed by two cliques K_3.

Define for a connected graph G the *distance* $d(u, v)$ between a pair of vertices u and v as the length of the shortest path from u to v.

Theorem 58 *A graph G is bipartite if and only if it does not contain a cycle of odd length.*

Proof: Let $G = (V, E)$ be a bipartite graph with the vertex set V partitioned into the independent sets X and Y. Consider a cycle $v_1 v_2 \cdots v_k v_1$. This vertex sequence has to alternate between X and Y implying that k must be even; e.g. $v_1 \in X, v_2 \in Y, \ldots, v_k \in Y, v_1 \in X$.

On the other hand, let $G = (V, E)$ be a graph without odd cycles. We can assume without loss of generality that G is connected. Let u be an arbitrary vertex in V. Split V into the two disjoint sets:

$$X = \{v \in V \mid d(u, v) \text{ is even}\}$$
$$Y = \{v \in V \mid d(u, v) \text{ is odd}\}$$

We claim that X and Y are the desired bi-partition of V. Namely that X and Y are independent sets. Clearly $X \cup Y = V$ and $X \cap Y = \emptyset$. Suppose on the contrary that two vertices z and w both in X or both in Y are connected by an edge. Consider the shortest paths $P_z = z_1, \ldots, z_k$ from $u = z_1$ to $z = z_k$ and $P_w = w_1, \ldots, w_l$ from $u = w_1$ to $w = w_l$. Let $u' = z_j = w_i$ be the last vertex both in P_w and in P_z. See figure 10.7. Since P_z and P_w were shortest paths, $j = i$. By the definition of X and Y, $k + l$ is even. Thus, the walk $u = z_j, \ldots, z_k = z, w = w_l, \ldots, w_i = u$ is an odd cycle, in contradiction to the assumption that G has no odd cycles. $\Box$

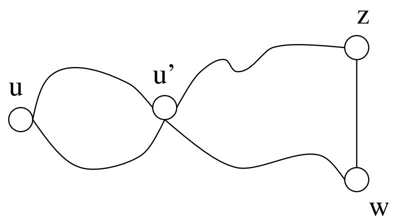

Figure 10.7:

Exercise 59 *Prove that the distance $d(u, v)$ is a metric. That is:*

1. $d(x, x) = 0$
2. $d(x, y) = d(y, x)$*; and*
3. $d(x, z) + d(z, y) \geq d(x, y)$.

Let $D(u, v)$ be the maximum length of a path connecting u and v. Prove that also $D(u, v)$ is a metric.

Chapter 11

Trees

In this section we will give some basic properties of trees. A *tree* is defined as a connected acyclic simple undirected graph.

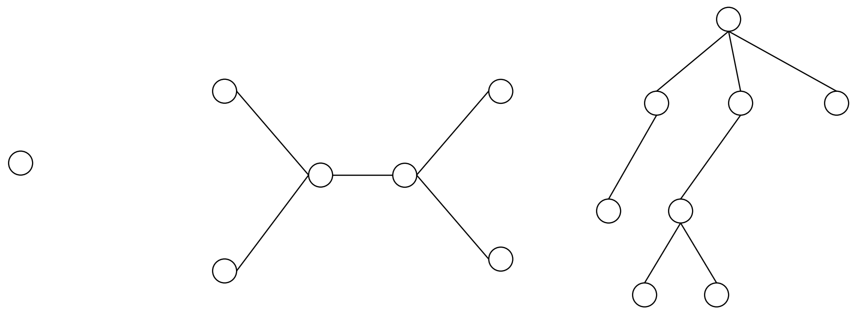

Figure 11.1: Some examples of trees

Theorem 60 *Let T be a tree and let u, v be vertices in T, $u \neq v$. Then, there is a unique path in T connecting u and v.*

Proof: There must be a path from u to v in T since T is connected by definition. Suppose that there are two different paths P_1 and P_2 connecting u and v. Assume that there is an edge $e = xy$ on the path P_1 and not in P_2. Consider the subgraph $H = (P_1 \cup P_2) - e$ containing all the edges and vertices on the paths P_1 and P_2, except the edge e. H has to be connected since u and v are connected by P_2 and any vertex on P_1 is either connected to u or v. Therefore, x and y are connected by a path $P_3 = x, \ldots, y$ in H. But then we have the cycle $P_3 + xy$ in T contradicting the assumption of T being a tree and thus acyclic. □

Theorem 61 *A tree T with n vertices has $n - 1$ edges.*

Proof: We prove the theorem by induction on the number of vertices n. The base for $n = 1$ is trivial. For the inductive step, assume that the claim holds for any tree with $n - 1 \geq 1$ vertices and let T be a tree with n vertices. Let P be the longest path in T and u be one of the end vertices of P. $d(u) = 1$ since we otherwise could extend the path to be even longer. Delete from T the vertex u and the edge incident to u, obtaining a tree T' with $n - 1$ vertices. By the induction hypothesis T' must have $n - 2$ edges and thus T has $n - 1$ edges. □

Exercise 62 *A sequence of integers $d_1, \ldots, d_n$ is the degree sequence of some tree if and only if:*

1. *$d_i \geq 1$ for $i = 1, \ldots, n$; and*
2. *$\sum_{i=1}^{n} d_i = 2n - 2$.*

Solution: The condition is necessary by Theorems 55 and 61. We show by induction on the number of vertices that the condition is also sufficient. The induction base for $n = 1$ and $n = 2$ is immediate. Assume that the claim holds for any sequence of $n - 1 \geq 2$ integers. Since $\sum d_i = 2n - 2$, there must be at least one $d_i = 1$ (there are actually at least two) and at least one $d_j \geq 2$. Removing d_i from the sequence and substructing one from d_j we obtain a new sequence that satisfies the inductive requirements. By the inductive hypothesis there exists a tree with the new degree sequence. Connect a new vertex to the vertex of degree $d_j - 1$ in this tree, obtaining a tree with degree sequence $d_1, \ldots, d_n$.

11.1 Cut-vertices and cut-edges

Given a graph $G = (V, E)$, let $\omega(G)$ be the number of connected components in G. A vertex $v \in V$ is called a *cut-vertex* if $\omega(G - v) > \omega(G)$. Similarly an edge $e \in E$ is called a *cut-edge* if $\omega(G - e) > \omega(G)$. Cut-vertices and cut-edges are simple concepts, but they have important uses in these lectures.

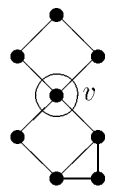

Figure 11.2: v is the only cut vertex in this graph.

We give next some simple observations about cut-vertices and cut-edges. The ommitted proofs are left as exercises.

Lemma 63 *Given a tree T, a vertex $v \in T$ is a cut-vertex if and only if $d(v) > 1$.*

Proof: If $d(v) = 0$ or $d(v) = 1$, then clearly v is not a cut vertex. Assume that $d(v) \geq 2$ and let x and y be two vertices adjacent to v. By Theorem 60, there is a unique path in T between x and y and so this path must be xv, vy, establishing that v is a cut vertex. The other direction is proved similarly. □

We will see soon that every graph has a spanning tree, so by the last lemma any non-trivial graph has at least two vertices that are not cut-vertices.

Lemma 64 *In a graph $G(E, V)$, $e \in E$ is a cut edge if and only if it is not in a cycle.*

Lemma 65 *Let G be a connected graph. G is a tree if and only if every edge in G is a cut edge.*

Corollary 66 *Every connected graph G contains a spanning tree.*

Proof: Start removing edges that are not cut edges until left with a *minimal* connected graph. Since every edge in the remaining graph is a cut edge it follows by Lemma 65 that we have got a tree. □

Corollary 67 *Every connected graph G has at least $|V| - 1$ edges.*

Proof: By Corollary 66, G contains a spanning tree having $|V| - 1$ edges. □

Corollary 68 *If we add a new edge e to a tree T, then the resulting graph $T + e$ has a unique cycle.*

Proof: The added edge is not a cut edge, so by Lemma 64 it is in a cycle. Since T is a tree, every cycle in $T + e$ must contain $e = (x, y)$. Suppose there are two cycles. Then there are two different paths in the tree T between the vertices x and y, in contradiction with Theorem 60. □

11.2 Cayley's formula

We next consider the problem of enumerating the labeled spanning trees of a connected graph $G = (V, E)$. Denote this number by $t(G)$. It is possible to count this for any graph, but we shall restrict ourselves to $G = K_n$ the complete graph. We will suggest three solutions for the problem, leaving some details as an exercise.

Before starting out, though, we mention that the solution for general graphs can be obtained by introducing the notion of *contraction* which means "collapsing an edge" and identifying the two vertices on its end points, leaving the other edges intact. The graph resulting from the contraction of an edge e in G is denoted $G \cdot e$.

The number of spanning trees can be counted by choosing an arbitrary edge e and counting separately the spanning trees that do not include the edge and those that include the edge. The spanning trees that do not include the edge e are exactly the spanning trees of $G - e$. The spanning trees that include the edge are those of the contracted graph $G \cdot e$. As an example, consider

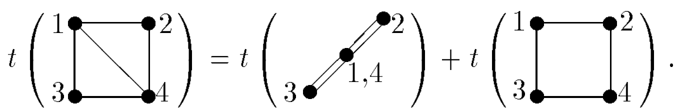

Observe that when we contract the edge, we must consider the mutiple edges the result from the contraction.

I–Generating Functions

Define $t_n := t(K_n)$. Then, $t_0 = 0$, $t_1 = t_2 = 1$, and from the previous chapters we have two possible candidates for the generating function:

$$\sum_n t_n z^n \qquad \text{or} \qquad \sum_n t_n \frac{z^n}{n!}.$$

The latter exponential generating function is particularly well-suited for our purposes, as we will discover. But first, we want to determine a recurrence relation for the t_n's.

The central idea for our choice of a recurrence will be this: choose in a spanning tree T of K_n a vertex v and remove all edges incident on v. To illustrate this, we give the following example:

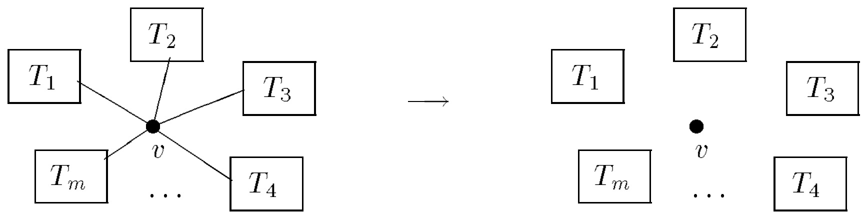

Since v is a vertex in a spanning tree, v will connect m subtrees $T_1, \ldots, T_m$ that have no vertices in common. The removal of all edges incident to v will therefore leave us standing with m blocks of different sizes. Assume in the following that the number of vertices in T_i equals k_i for $i = 1, 2, \ldots, m$.

It is time for a little definition:

Definition 69 (Multinomial coefficient) *The multinomial coefficient is defined as*

$$\binom{n}{k_1, \ldots, k_m} := \frac{n!}{k_1! k_2! \ldots k_m!}.$$

The mutinomial coefficients give then number of ways to partition $n = k_1 + \cdots + k_m$ items into m classes of sizes $k_1, \ldots, k_m$:

$$\begin{aligned}
\binom{n}{k_1, \ldots, k_m} &= \binom{n}{k_1}\binom{n-k_1}{k_2}\cdots\binom{n-k_1-\cdots-k_{m-1}}{k_m} \\
&= \frac{n!}{k_1!(n-k_1)!}\frac{(n-k_1)!}{k_2!(n-k_1-k_2)!}\cdots\frac{(n-k_1-\cdots-k_{m-1})!}{k_m!(n-k_1-\cdots-k_m)!} \\
&= \frac{n!}{k_1!k_2!\cdots k_m!}.
\end{aligned}$$

To avoid triviality, we may assume $n > 1$ and $m > 0$. Now consider the complete subgraphs $K_{k_1}, \ldots, K_{k_m}$. The product of the number of spanning trees of the subgraphs is $t_{k_1} t_{k_2} \ldots t_{k_m}$, and the number of ways of combining them is $k_1 \cdot k_2 \cdots k_m$. From this we get the desired recurrence for t_n (observe that $k_i \geq 1$):

$$t_n = \sum_{m>0} \frac{1}{m!} \sum_{k_1+\cdots+k_m=n-1} \binom{n-1}{k_1, \ldots, k_m} t_{k_1} \cdot t_{k_2} \cdots t_{k_m} k_1 \cdot k_2 \cdots k_m,$$

where we have divided by $m!$ since the order of the classes (T_i's) is of no importance. We further simplify this expression by introducing numbers u_n, $u_n := t_n \cdot n$:

$$\begin{aligned}\frac{u_n}{n} &= \sum_{m>0} \frac{1}{m!} \sum_{k_1+\cdots+k_m=n-1} \binom{n-1}{k_1,\ldots,k_m} u_{k_1} \cdot u_{k_2} \cdots u_{k_m} \\ &= \sum_{m>0} \frac{1}{m!} \sum_{k_1+\cdots+k_m=n-1} \frac{(n-1)!}{k_1!k_2!\cdots k_m!} u_{k_1} \cdot u_{k_2} \cdots u_{k_m}.\end{aligned}$$

Thus,

$$\frac{u_n}{n!} = \sum_{m>0} \frac{1}{m!} \sum_{k_1+\cdots+k_m=n-1} \frac{u_{k_1}}{k_1!} \frac{u_{k_2}}{k_2!} \cdots \frac{u_{k_m}}{k_m!},$$

and we may form the exponential generating function of the series u_n, $n \in \mathsf{N}$:

$$\sum_{n>1} \frac{u_n}{n!} z^n = \sum_{n>1} \left(\sum_{m>0} \frac{1}{m!} \sum_{k_1+\cdots+k_m=n-1} n! \frac{u_{k_1}}{k_1!} \frac{u_{k_2}}{k_2!} \cdots \frac{u_{k_m}}{k_m!} \right) \frac{z^n}{n!}. \tag{11.1}$$

Recall that for $\hat{A}(z) := \sum_n a_n \frac{z^n}{n!}$ and $\hat{B}(z) := \sum_n b_n \frac{z^n}{n!}$:

$$\hat{A}(z)\hat{B}(z) = \sum_n \sum_r \binom{n}{r} a_r b_{n-r} \frac{z^n}{n!} = \sum_n \sum_{r+s=n} \binom{n}{r} a_r b_s \frac{z^n}{n!}.$$

This can be generalised for more than two functions. If we take

$$\begin{aligned}\hat{A}_1(z) &:= \textstyle\sum_n a_{1,n} \frac{z^n}{n!}, \\ \hat{A}_2(z) &:= \textstyle\sum_n a_{2,n} \frac{z^n}{n!}, \\ &\vdots \\ \hat{A}_m(z) &:= \textstyle\sum_n a_{m,n} \frac{z^n}{n!},\end{aligned}$$

then

$$\begin{aligned}&\hat{A}_1(z)\hat{A}_2(z)\cdots\hat{A}_m(z) = \\ &\qquad \sum_n \sum_{r_1+\cdots+r_m=n} \binom{n}{r_1,\ldots,r_m} a_{1,r_1} a_{2,r_2} \cdots a_{m,r_m} \frac{z^n}{n!}.\end{aligned}$$

With this in mind, we define $\hat{U}(z) := \sum_n u_n \frac{z^n}{n!}$ and recall equation (11.1):

$$\begin{aligned}\sum_{n>1} \frac{u_n}{n!} z^n &= \sum_{n>1} \left(\sum_{m>0} \frac{1}{m!} \sum_{k_1+\cdots+k_m=n-1} n! \frac{u_{k_1}}{k_1!} \frac{u_{k_2}}{k_2!} \cdots \frac{u_{k_m}}{k_m!} \right) \frac{z^n}{n!} \\ &= \sum_{n>1} n \sum_{m>0} \frac{1}{m!} \left\langle \frac{z^{n-1}}{(n-1)!} \right\rangle (\hat{U}(z))^m \frac{z^n}{n!}\end{aligned}$$

$$
\begin{aligned}
&= \sum_{n>1}\left(n\left\langle\frac{z^{n-1}}{(n-1)!}\right\rangle\sum_{m>0}\frac{1}{m!}(\hat{U}(z))^m\right)\frac{z^n}{n!}\\
&= \sum_{n>1}\left\langle\frac{z^{n-1}}{(n-1)!}\right\rangle z(e^{\hat{U}(z)}-1)\frac{z^{n-1}}{(n-1)!}\\
&= \sum_{n>0}\left(\left\langle\frac{z^n}{n!}\right\rangle z(e^{\hat{U}(z)}-1)\right)\frac{z^n}{n!}\\
&= z(e^{\hat{U}(z)}-1).
\end{aligned}
$$

Thus, $\hat{U}(z) = ze^{\hat{U}(z)}$. This expresses $\hat{U}$ in terms of itself, and we will (referring to complex analysis) just have to accept that a solution of this equation is given by

$$\hat{U}(z) = z\hat{E}(z), \qquad \hat{E}(z) = \sum_k (k+1)^{k-1}\frac{z^k}{k!}.$$

We may now calculate u_n as

$$
\begin{aligned}
u_n &= \left\langle\frac{z^n}{n!}\right\rangle \hat{U}(z) = \left\langle\frac{z^n}{n!}\right\rangle z\hat{E}(z)\\
&= n!\left\langle z^{n-1}\right\rangle \hat{E}(z) = n!\frac{n^{n-1}}{(n-2)!}\\
&= n\cdot n^{n-2}.
\end{aligned}
$$

Finally, since $t_n = \frac{u_n}{n}$, we arrive at

$$\textit{Cayley's formula:} \qquad t_n = n^{n-2}.$$

II—The Prüfer Code

Cayley's formula can be obtained by much simpler methods and without the heavier machinery we have used above.

Suppose A and B are two sets. If $f : A \to B$ is an injective function (1–1-mapping), this means that two distinct elements of A cannot have the same image under f; therefore we must have that $|B| \geq |A|$. If there also exist an injective function from B to A, this combines to $|A| = |B|$. We will make use of this fact and establish a 1–1 correspondence between spanning trees and certain numeral encodings of them, called *Prüfer codes*.

Definition 70 (Prüfer code) *Assume that the nodes V of a tree T are indexed, such that $V = [n]$. To obtain the Prüfer code of T, repeat the following step until all edges have been removed:*

> *Pick among the nodes with degree 1 the node v with least index. Delete both v and the edge $\{v, w\}$ from T, and append w to the Prüfer code.*

At the end remove the last number from the Prüfer code (as it will show up, this number is spurious).

The following fact can be easily proved by induction.

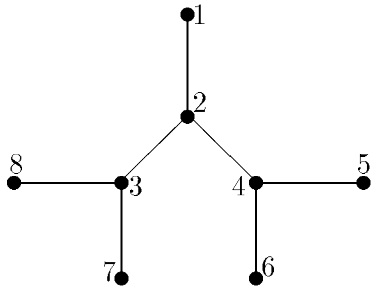

The Prüfer of the (spanning) tree above is:

$$\overbrace{[244233]}^{n-2}.$$

Figure 11.3: An example of a spanning tree and its Prüfer code.

Fact 71 *In the Prüfer code of a tree, each vertex v appears exactly $d(v) - 1$ times.*

Proof: One may prove this fact by induction on n, the case 1 being obvious. For the induction step, consider the first step where the leaf v and the edge $\{v, w\}$ have been removed from T; the code thus begins $[w \ldots$, and we are left with a subtree T'. By induction, any vertex $u \in V \setminus \{v, w\}$ appears $d_{T'}(u) - 1 = d_T(u) - 1$ times. Now, v is a leaf, so $d_T(v) - 1 = 0$; finally, $d_{T'}(w) - 1 = d_T(w) - 2$, and we note that w has appeared already once in the code for T. □

Now, let A be the set of all spanning trees of the complete graph K_n, and let B be the set $\{[a_1, \ldots, a_{n-2}] \mid a_i \in [n]\}$ of all $n - 2$ tuples of integers from $[n]$. Then , $|B| = n^{n-2}$, and the Prüfer encoding is a mapping from A to B.

The reverse mapping is similarly defined. For the purpose of counting it suffices to show that the mapping is injective. Consider the following mapping from B to A (i.e. the construction of a tree), given any sequence $[a_1, \ldots, a_{n-2}]$:

- Let $a_{n-1} := n$.
- For $i = 1, 2, \ldots, n - 1$, let b_i be the least number not included in $\{a_i, a_{i+1}, \ldots, a_{n-1}\} \cup \{b_1, \ldots, b_{i-1}\}$.
- Then, $T = (V, E)$, $V := [n]$, $E := \{\{b_i, a_i\} \mid i = 1, 2, \ldots, n - 1\}$ will be a spanning tree of K_n.

We now claim that that both mappings are 1–1, and that by this $|A| = |B| = n^{n-2}$. Let us first examine the Prüfer encoding:

Lemma 72 *Let $T \neq T'$ be two trees. Then $P(T) \neq P(T')$.*

Proof: By induction on the number of nodes of the tree:

If the trees differ in the least leaf, surely the lemma holds (some vertex appears in $P(T)$ but not in $P(T')$.). Secondly, if the least label is the same in T and T' but is attached to different vertices, then the first elements of $P(T)$ and $P(T')$ will be different as well, and the lemma holds again.

The rest easily follows by induction: suppose the least leaf v is attached to vertex u. Then (with sloppy notation) we only need to observe that $P(T) = [u, P(T \setminus \{u, v\})]$ and $P(T') = [u, P(T' \setminus \{u, v\})]$. □

Exercise 73 *Show that the mapping above from B to A is 1–1 and that it is the inverse of the mapping from A to B. (Hint: it may be useful to argue by contradiction.)*

Exercise 74 *Prove that the number of spanning trees of the complete graph K_n with degrees $d_1, \ldots, d_n$ is $\binom{n-2}{d_1-1, d_2-1, \ldots, d_n-1}$.*

III–Direct counting

We leave most part of this solution as an exercise and only outline the steps that the proof follows.

1. Prove that the number of spanning trees of the complete graph K_n with degree sequence $d_1, \ldots, d_n$ is $\binom{n-2}{d_1-1, d_2-1, \ldots, d_n-1}$. This can be proved directly by extending the solution of Exercise 62.

2. Recall the binomial theorem $(x + y)^n = \sum \binom{n}{r} x^r y^{n-r}$. There is a similar *multinomial theorem*:

$$(x_1 + \cdots + x_m)^n = \sum_{r_1 + \cdots + r_m = n} \binom{n}{r_1, \ldots, r_m} x_1^{r_1} x_2^{r_2} \cdots x_m^{r_m}.$$

   Count the number of trees by taking the sum of the number of trees with any degree sequence:

$$(1 + 1 + \cdots + 1)^{n-2} = n^{n-2} = \sum_{d_1' + \cdots + d_n' = n-2} \binom{n-2}{d_1', \ldots, d_n'}.$$

11.3 Application: lower bounds

We will now give some simple lower bounds in the comparison decision tree model.

Theorem 75 *Any comparison decision tree algorithm to find the maximum of n elements requires at least $n - 1$ comparisons.*

Proof: Let G be a graph with n vertices, one for each input element. Initially G has no edges. Every time the algorithm compares two elements we add an edge between the corresponding vertices. Clearly, as long as the graph is not connected, the algorithm can not establish the relative order between some of the input elements. We have seen that for a graph to be connected it must have at least $n - 1$ edges, giving the lower bound. □

The following lemma will be used in the sorting lower bounds given next.

Lemma 76 *A tree (directed from parent to children) with outdegree at most d and height h has at most d^h leaves.*

Proof: By induction on h. If $h = 0$, then any tree of height h has $d^0 = 1$ vertex. Assume that the claim holds for any tree with height at most $h - 1$. Let T be a tree with height h. Removing the root from T leaves us with at most d subtrees of height at most $h - 1$. By the induction hypothesis the number of leaves in each such subtree is at most d^{h-1} and the number of leaves in T is at most $d \cdot d^{h-1} = d^h$. □

Theorem 77 *Sorting n distinct elements requires at least $\Omega(n \log n)$ comparisons.*

Proof: Sorting n elements is a question of finding out which permutation of the sorted list the input is. Since there are $n!$ distinct permutations, we have $n!$ distinct outcomes. So there must be at least $n!$ leaves in the comparison decision tree. Since each node in the tree corresponds to a comparison with two possible outcomes, the comparison decision tree has outdegree 2. So we can compute the promised lower bound for the height h of the tree by:

$$\begin{array}{lrcl} & 2^h & \geq & n! \\ \Downarrow & & & \\ & h & \geq & \log n! \\ \Downarrow & & & \\ & h & = & \Omega(n \log n) \quad \text{either by Stirling's formula or since } n! \geq (\frac{n}{2})^{\frac{n}{2}}. \quad \square \end{array}$$

But what if we instead of one comparison in each step could make p comparisons using p processors in parallel?

This model is called Valiant's parallel comparison decision tree model.

The following lemma will used in the parallel sorting lower bound given next.

Lemma 78 *Suppose G is an n-vertex graph with p edges. Then the number of acyclic orientations of G's edges, denoted $aco(G)$, is at most $(1 + \frac{2p}{n})^n$.*

Proof: Let $G = (V, E)$ be an undirected graph with p edges and n vertices. If d_i is the degree of vertex v_i, by Lemma 55, $\sum_{i=1}^{n} d_i = 2p$. For each vertex, we have $d_i + 1$ ways to choose the number of outgoing edges, i.e. the outdegree d_i^+. This gives $\Pi_{i=1}^{n}(d_i + 1)$ ways to choose the outdegrees d_i^+ for all vertices. Clearly every acyclic orientation determines the outdegrees of all vertices. However not every choice of outdegrees corresponds to a single acyclic orientation. If fact, some choices do not correspond to any orientation: for example $d_i^+ = d_i$ for all vertices v_i. Suppose that some choice of outdegrees corresponds to an acyclic orientation. We argue that then the acyclic orientation is unique. We can show this by induction on n. For $n = 1$ it holds vacuously, since we cannot have any edges. For $n > 1$ we must have (since n is finite and we have no cycles) a vertex v_i with $d_i^+ = 0$. Pick that vertex and orient all it's edges as in-coming. All orientations must agree on this choice. By induction they also agree on the graph which remains if we remove v_i and it's in-coming edges. The conclusion is, that $\Pi_{i=1}^{n}(d_i + 1)$ is an upper bound for $aco(G)$. By using the well known relation between the geometric and the arithmetic mean (geometric mean is always smaller than arithmetic mean), we get

$$\begin{aligned} aco(G) &\leq \Pi_{i=1}^{n}(1+d_i) \\ &= (\sqrt[n]{\Pi_{i=1}^{n}(1+d_i)}\,)^n \\ &\leq (\frac{\sum_{i=1}^{n}(1+d_i)}{n})^n \\ &= (1+\frac{2p}{n})^n. \quad \square \end{aligned}$$

Theorem 79 *Sorting n distinct elements with p processors requires at least*

$$\Omega(\frac{n\log n}{p}+\frac{\log n}{\log(1+\frac{2p}{n})})$$

rounds.

Proof: At a first glance, it seems that there are 2^p possible outcomes at each step, i.e. the nodes in our comparison tree have outdegree 2^p. This gives the following lower bound on the height:

$$\begin{aligned} (2^p)^h &\geq n! \\ \Downarrow \quad & \\ h &\geq \tfrac{\log n!}{\log 2^p} \\ \Downarrow \quad & \\ h &= \Omega(\tfrac{n\log n}{p}). \end{aligned}$$

This is just the trivial lower bound we could get from the one processor lower bound, by arguing that 1 processor can simulate p processors in time O(p).

Can we do better? Yes. The main observation is that not all 2^p answers to the comparisons make sense. To illustrate this suppose we have $\binom{n}{2}$ processors, i.e. one processor for each possible pair of elements. From a first glance, there are $2^{\binom{n}{2}}$ different outcomes, but we have only $n! << 2^{\binom{n}{2}}$ different permutations of the n elements, so a lot of outcomes do not correspond to permutations of the input. The key is that answers like $a_1 < a_2 < a_3 < a_1$ are not possible when $<$ is an transitive ordering relation. To be more general: The numbers of answers from p comparisons, in a set of n elements, is bounded by the number of ways you can orient p edges in an n-vertex graph without obtaining a cycle. By Lemma 78, the number of such acyclic orientations is bounded by $(1+\frac{2p}{n})^n$. This gives us a new lower bound on the height of the parallel comparison decision tree:

$$\begin{aligned} ((1+\tfrac{2p}{n})^n)^h &\geq n! \\ \Downarrow \quad & \\ h &\geq \tfrac{\log n!}{n\log(1+\frac{2p}{n})} \\ \Downarrow \quad & \\ h &= \Omega(\tfrac{n\log n}{n\log(1+\frac{2p}{n})}) = \Omega(\tfrac{\log n}{\log(1+\frac{2p}{n})}). \quad \square \end{aligned}$$

The lower bound proven above for parallel sorting is actually tight as there exist comparison model sorting algorithms that achieve this bound (we will not see such algorithms however in this course).

Chapter 12

Eulerian and Hamiltonian walks

Eulerian and Hamiltonian walks are walks that cover all edges and vertices in a graph.

12.1 Euler tours

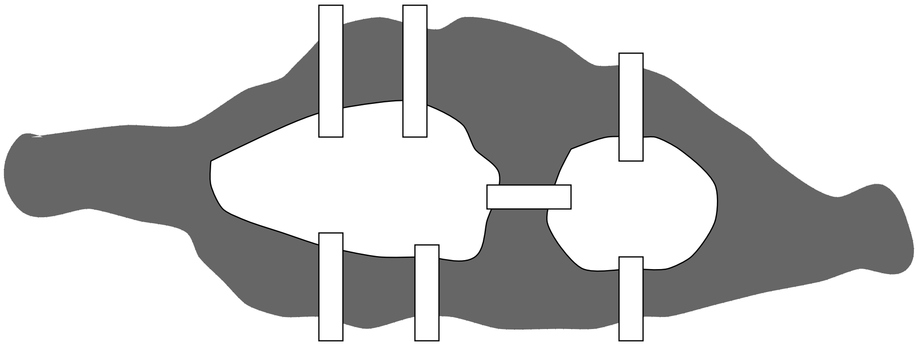

Figure 12.1: City map of Königsberg

In 1736 L. Euler wrote a paper, which is considered the first paper in Graph Theory, describing the problem of the 7 bridges of Königsberg (Figure 12.1). The problem is of interest for tourists and sounded like this: Is it possible to take a walk starting and ending in the same point, crossing every bridge exactly once?

This leads to the following definition: An *Euler circuit* in G is a closed walk that traverses each edge of G exactly once. G is called *Eulerian* if it has an Euler circuit. An *Euler trial* in G is a trial that traverses each edge of G exactly once.

To solve the problem of the city walk, we will have a criteria for when a graph is Eulerian.

Theorem 80 *A connected graph G has an Euler circuit if and only if all vertices have even degree.*

Proof: That the condition is necessary should be obvious. To show it is sufficient, we show how to construct the circuit by forming edge disjoint closed circuits and patching them together. Pick an arbitrary vertex and begin traversing an arbitrary path from it, using only edges that have not been traversed yet. Since every vertex has even degree, the walk will always end at the starting vertex. If the walk does not cover all the edges, we can find a vertex on circuit with edges that have not been used yet, (since the graph is connected) and begin forming a new circuit. These two circuit can be merged into one large circuit and we can proceed like this until we have included all edges. □

Studying the map (The ground is white, and the water black), we can see that no vertices (island/shores) have even degree, so the answer must be that there is no tour that covers each bridge precisely once.

If we should plan a tour in a car in Aarhus C, it would be necessary to consider directed edges. It is easy to show with a similar proof, that the sufficient and necessary condition for a directed connected graph to be Eulerian, is that the out-degree equals the in-degree for every vertex.

We have the following criteria for when we can find an Euler trail:

Corollary 81 *A connected graph G has an Euler trail if and only if it has exactly two vertices with odd degree.*

Proof: Suppose we have exactly two vertices of odd degree. If we could connect these two, we would have an Euler circuit, so we have an Euler path going from one of the edges with odd degree, and ending in the other. Conversely, if we have an Euler trail, it is easy to see that it must begin and end in vertices with odd degree, and that all other vertices have even degree. □

12.2 Hamilton cycles

A *Hamiltonian cycle* in G is a cycle that passes every vertex of G exactly once. A *Hamiltonian path* in G is a path that visits every vertex in G exactly once.

One problem with Hamiltonian cycles is that it is very difficult to decide whether a graph has a Hamiltonian cycle. As an example take a look at the two graphs in Figure 12.2 and try to find a Hamiltonian cycle in them.

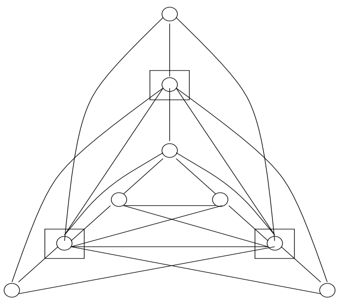

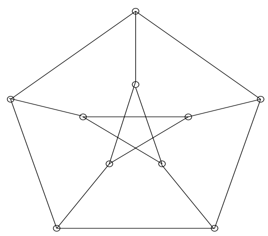

Figure 12.2: Two non-Hamiltonian graphs.

Theorem 82 *If G is a Hamiltonian graph, then for any subset of vertices S, $\emptyset \neq S \subseteq V$ we have that $\omega(G \backslash S) \leq |S|$, where $\omega(G \backslash S)$ denotes the number of connected components in G.*

Proof: Assume we have a Hamiltonian graph $G = (V, E)$. Then we have a Hamiltonian cycle C. Let us see what happens to the connectivity when we remove a non-empty set of vertices S from the cycle.

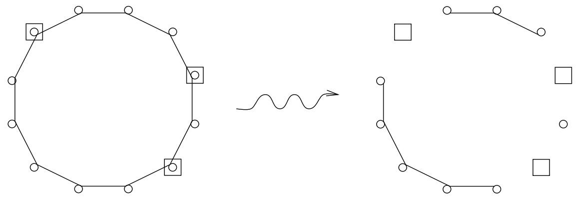

Figure 12.3: Removing vertices form a cycle.

From Figure 12.3 it is clear that the cycle breaks into at most S parts, and thus, $\omega(C\backslash S) \leq |S|$. As the Hamiltonian cycle is a subgraph of G, considering the whole graph we get that $\omega(G\backslash S) \leq \omega(C\backslash S)$, completing the proof. □

The graph on the left in Figure 12.2 is non-Hamiltonian since by the last theorem one can see that removing the 3 marked vertices gives a graph with 4 connected components.

Exercise 83 *Show that Petersen's graph is not Hamiltonian. (Hint: Do not use the theorems from this chapter.)*

When we consider Hamiltonian graphs we do not care whether the graphs have self loops or multiple edges. A Hamiltonian cycle visits each vertex exactly once and therefore it is not a limitation to look at simple graphs only. (You may add all the loops you want and the graph will not change status with respect to Hamiltonian cycles). The next theorem of Dirac can be used to prove that a graph is Hamiltonian by looking only at the degree of the vertices in the graph. (By the way, Dirac was a professor at Aarhus University!)

Theorem 84 (Dirac) *Let $G = (V, E)$ be a simple graph with $n = |V| \geq 3$. If $\delta(G) \geq n/2$, then G is Hamiltonian.*

Proof: The proof proceeds by considering a maximal counterexample G – maximal in the sense that we add edges to the graph as long as we do not introduce a Hamiltonian cycle.

That K_n the complete graph is Hamiltonian is clear. Assume that we have a simple graph G with $|V| \geq 3$ and $\delta \geq n/2$ which is not Hamiltonian, but $G + (u, v)$ is. Then G has a Hamiltonian path $u = v_1, \ldots, v_n = v$. See Figure 12.4

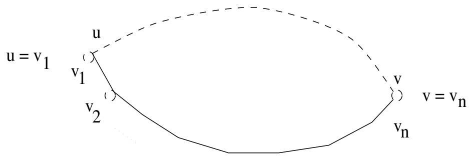

Figure 12.4: A Hamiltonian cycle in $G + uv$ gives a Hamiltonian path in G.

Now we define two sets S and T

$$S = \{v_i \mid uv_{i+1} \in E\} \qquad T = \{v_i \mid v_i v \in E\}$$

and claim that

$$|S \cup T| \leq n - 1$$

as $v_n = v$ is not in S and neither is in T. Now we claim that

$$S \cap T = \emptyset$$

for otherwise if $v_i \in S \cap T$, then G has the Hamiltonian cycle $v_1, \ldots, v_i, v_n, \ldots, v_{i+1}, v_1$, in contradiction with the assumption that G is not Hamiltonian. See Figure 12.5.

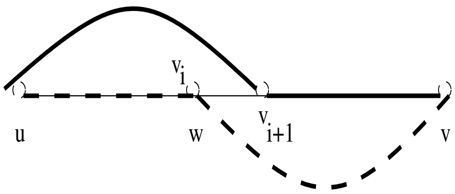

Figure 12.5: A Hamiltonian cycle in G.

Now we notice, that the degree of u and v is given by the cardinality of the sets S and T

$$d(u) = |S| \qquad\qquad d(v) = |T|$$

and we get:

$$\begin{aligned} d(u) + d(v) &= |S| + |T| \\ &= |S \cup T| \quad \leq \quad n - 1 \end{aligned}$$

in contradiction with the assumption that $d(u), d(v) \geq n/2$. This completes the proof. □

Similarly one can obtain the following lemma.

Lemma 85 *Let $G = (V, E)$ be a simple graph with an edge $(u, v) \notin E$, such that $d(u)+d(v) \geq n$. Then G is Hamiltonian if and only if $G + (u, v)$ is Hamiltonian.*

Proof: If G is Hamiltonian, then we may add all the edges we like and the Hamiltonian cycle is still there. The same proof by contradiction as in Dirac's theorem gives the other direction of the proof. □

The last lemma can be applied successively, motivating the following definition.

Definition 86 (The closure of a graph) *Given a simple graph $G = (V, E)$, the closure of the graph $C(G)$ is defined as the graph obtained by repeatedly choosing all pairs of vertices $u, v \in V$ which satisfies $(u, v) \notin E$ and $d(u) + d(v) \geq n$, and adding edges between them.*

But is this graph well defined? We show that the order in which we pick the vertices to be connected is irrelevant for the graph $C(G)$.

Lemma 87 *$C(G)$ is well defined. That is $C(G)$ obtained by adding the edges*

$$e_1 e_2 \cdots e_k$$

yields the same graph as the one obtained by adding the edges

$$f_1 f_2 \cdots f_l.$$

Proof: The proof is by contradiction. Assume that the two sequences above correspond to two different closures of a simple graph $G = (V, E)$. Now take the edges from the e-sequence that also appear in the f-sequences, and stop the first time you find an edge in the e-sequence which is not in the f-sequence, that gives us

$$\{e_1, \cdots, e_{k'}\} \subseteq \{f_i\} \quad k' \leq k.$$

and

$$e_{k'+1} \notin \{f_i\}.$$

Now let $H = G + e_1 + \cdots + e_{k'}$ and let $e_{k'+1} = (u, v),\ u, v \in V$. Then as a consequence of the definition of the closure we have

$$d_H(u) + d_H(v) \geq n.$$

But H is also a subgraph of $H' = G + f_1 + \cdots + f_l$ and therefore the inequallity also holds in H'

$$d_{H'}(u) + d_{H'}(v) \geq n.$$

But then we have also added the edge in the second closure graph which is a contradiction with $e_{k'+1} \notin \{f_i\}$. Hence the closure of a graph is well defined. □

The last lemma can be used to prove the following theorem which is a generalisation of Dirac's theorem.

Theorem 88 *Let $G = (V, E)$ be a simple graph. If $C(G) \cong K_n$, then G is Hamiltonian.*

Proof: The complete graph K_n is clearly Hamiltonian. By the lemma above if the closure of G is K_n, then G is also Hamiltonian. □

As a last criteria for a graph to be Hamiltonian we introduce the following theorem.

Theorem 89 *Let G be a simple graph with n vertices having degrees*

$$d_1 \leq d_2 \leq \cdots \leq d_n.$$

Suppose that there is no $k < n/2$ such that

$$d_k \leq k \quad \textit{and} \quad d_{n-k} < n - k,$$

then G is Hamiltonian.

Proof: Given a simple graph $G = (V, E)$, we show that under the conditions above the closure $C(G)$ is the complete graph K_n. Thus, if the closure $C(G)$ is indeed the complete graph then G has a Hamiltonian cycle.

The proof proceeds by contradiction. Assume now that $C(G)$ is not the complete graph. Then there exist $u, v \in V$ and $(u, v) \notin C(G)$ (otherwise $C(G) = K_n$). Among all the possible choices for such u and v, choose that pair u and v in such a way that

$$d'(u) + d'(v) \leq n - 1$$

is maximal, where d' is the degree in the closure graph. The upper bound on the sum of degrees is due to the fact that the edge (u, v) is not in the closure graph. Now define the two sets S and T to be

$$S = \{w \mid wv \notin E, w \neq v\} \quad T = \{w \mid wu \notin E, w \neq u\}.$$

Because we chose u and v so the sum of their degrees in the closure graph is maximal we know the following about the degree of the vertices in S and T

$$\begin{array}{ll} s \in S & d'(s) \leq d'(u) \\ t \in T & d'(t) \leq d'(v) \end{array}$$

for otherwise we should have picked u and v differently. The size of the set S can be determined to be

$$|S| = |V| - |\{v\}| - |\{w \mid wv \in E\}| = n - 1 - d'(v)$$

and in a similar way we get the size of T:

$$|T| = n - 1 - d'(u).$$

Without loss of generality we can assume that

$$d'(u) \leq d'(v)$$

as the degree of the one vertice has to be less than or equal to the degree of the other. This gives us that

$$|S| = n - 1 - d'(v) \geq n - 1 - (n - 1 - d'(u)) = d'(u) = k$$

and there are at least k vertices in the set S. Because the degree (in the closure graph) of the vertices in S is bounded by the degree of u, we get that the closure graph has at least k vertices with degree less than or equal to $d'(u) = k$.

To find the number of vertices in T, we have to work a little harder. Immediately we have

$$|T| = n - 1 - d'(u) = n - k - 1.$$

That is, there are $n - k - 1$ vertices in T and their degrees are bounded by the degree of v. More precisely we get that

$$d'(t) \leq d'(v) \leq n - 1 - d'(u) = n - 1 - k < n - k.$$

But we wanted $n - k$ vertices with degree less than $n - k$. Luckily we have chosen u and v in such a way that u is not in T (otherwise u would have a self loop and G is simple) and $d'(u) \leq d'(v)$ so we do have $n - k$ vertices with degree less than $n - k$. Now we have found a $k < n/2$ such that

$$d'_k \leq k \quad \text{and} \quad d'_{n-k} < n - k,$$

where $d'_1 \leq \cdots \leq d'_n$ was the degree sequence of the closure graph $C(G)$. Since the degrees in the closure graph are not smaller than the degrees in the original graph, we have also that:

$$d_k \leq k \quad \text{and} \quad d_{n-k} < n - k.$$

But our assumption was that there was no such k leading to a contradiction. That gives us that $C(G)$ is the complete graph and therefore G is Hamiltonian. □

12.3 De Bruijn Sequences

An example of an application of Euler circuits is that of De Bruijn sequences, which discuss in this section.

Given a set of q symbols $\{0, 1, \cdots, q-1\}$ there are clearly q^k different blocks of k of symbols. A sequence of l symbols has $l-k+1$ (overlapping) sub-blocks of length k, and therefore at most $l-k+1$ distinct sub-blocks.

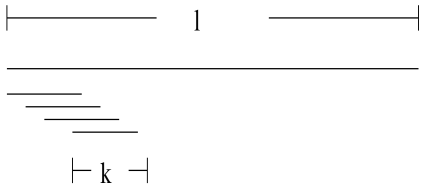

Figure 12.6: Sub-blocks of length k.

The question we are concerned with next is the construction of the shortest possible sequence that contains as sub-blocks all the distinct q^k k-blocks. Denote by $L(q,k)$ the length of such a sequence. It is clear that

$$L(q,k) \leq kq^k$$

as the sequence formed by taking all the distinct k blocks without overlap has this length. Since a sequence of length l has at most $l-k+1$ different sub-blocks of length k we get that

$$\begin{array}{lrcl} & q^k & \leq & L(q,k) - k + 1 \\ \Updownarrow & & & \\ & q^k + k - 1 & \leq & L(q,k) \\ \Downarrow & & & \\ & q^k + k - 1 & \leq & L(q,k) \leq kq^k \end{array}$$

But how short can we actually get the sequence? By formulating the problem as a problem on graphs we show that there exist sequences of the smallest possible length $L(q,k) = q^k + k - 1$. These sequences are called de Bruijn sequences.

Consider a circular sequence where the last sub-blocks are found by crossing the point where we start.

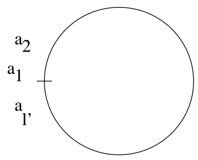

Figure 12.7: A circular sequence.

We construct a graph whose q^k vertices are labeled with all the different blocks of length k. Directed edges are connecting the vertex labeled

$$a_1 a_2 \cdots a_k.$$

to the vertices

$$a_2a_3\cdots a_{k-1}0 \quad a_2a_3\cdots a_{k-1}1 \quad \cdots \quad a_2a_3\cdots a_{k-1}[q-1].$$

These edges are labeled with the last symbol of the vertex label they lead to. Clearly, the circular sequence formed by the edge labels on a cycle that passes through all vertices (not necessarily all edges) will contain all the blocks of length k. If the cycle passes only one time in each node we get an optimal length sequence.

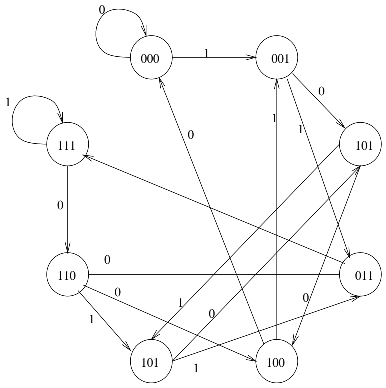

Figure 12.8: The graph with nodes labeled with length 3 blocks. The (cyclic) de Bruijn sequence starting at the vertex 000 is is 11101000.

Cycles that visit every node in a graph precisely once are Hamiltonian cycles. However, as we have seen, it is quite difficult to prove that a graph is Hamiltonian. Fortunately, there is another way to construct some graph so that the sequence we are looking for will form an Euler circuit, which are easier to deal with. In the new graph, the q^{k-1} vertices are labeled with blocks of length $k-1$ and the number of directed edges is q^k. Two vertices are connected if they are identical on $k-1$ positions:

$$a_1a_2a_3\cdots a_{k-1}$$

is connected by an edge directed to

$$a_2a_3\cdots a_{k-1}a_k.$$

And the label on the edge is

$$a_1a_2a_3\cdots a_{k-1}a_k.$$

As there are q^{k-1} vertices and each vertice has out-degree q (there are q strings which only differ on the last element) it is clear that the number of edges is $q^{k-1}q = q^k$ and that each edge is labeled with a different string.

If the graph would have an Eulerian circuit, it corresponds to a circular string that is formed by writing down the first symbol of each edge label. Clearly such a sequence would include all possible blocks of length k and it will have optimal length q^{k+1} (the sequence length which is also the length of the circular string). If we can argue that the graph is an Eulerian graph we immediately get that there exist a de Bruijn sequence of optimal length.

To see that the graph is an Eulerian graph remember how the edges were defined. Consider the following two vertices:

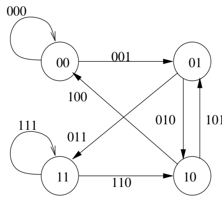

Figure 12.9: The graph with nodes labeled with length 2 blocks. Now edges are labeled with all the length 3 blocks and an Eulerian circuit corresponds to de Bruijn sequence.

$$\begin{array}{lllll} A = & a_1 & a_2 \cdots a_{k-2} & a_{k-1} & \\ B = & & a_2 \cdots a_{k-2} & a_{k-1} & a_k \end{array}$$

It is clear that there are q possible choices for a_k when $a_1, \ldots, a_{k-1}$ are fixed and that there are q choices for a_1 when $a_2, \ldots, a_k$ are fixed. Therefore, we have that $in{-}degree(u) = out{-}degree(u) = q$ for any vertex u, precisely the necessary and sufficient condition we have seen to have an Euler circuit. Hence there exists a de Bruijn sequence of optimal length $q^k + k - 1$.

Given an undirected graph $G = (V, E)$ we define its line graph $G^* = (V', E')$ as follows: $V' = E$ and for any two adjacent edges e_1 and e_2 in E (in V') there an edge (e_1, e_2) in E'.

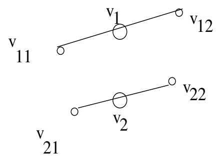

Figure 12.10: Two edges are transformed to two vertices

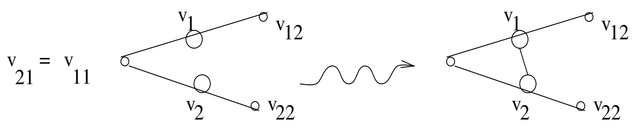

Figure 12.11: Two vertices in the line graph are connected.

The definition for a line graph of a directed graph is the same as for the undirected graphs only in the directed case we have to take care of the direction of the edges. That is, we have an edge directed form u to v in the directed line graph if in the original graph we have an edge corresponding to u that is directed into some vertex x and an edge corresponding to v that is directed away from the vertex x. See Figure 12.12.

The definitions lead us to state the following connection between Eulerian and Hamiltonian cycles.

Fact 90 *If we have an Euler circuit in a graph then we have a Hamiltonian cycle in its line graph. This is not necessarily true the other way around. (Observe also that not every graph is a line graph of another graph.)*

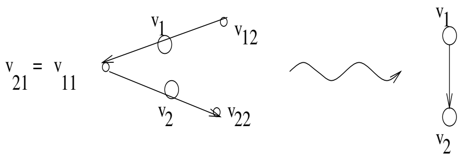

Figure 12.12: Directed line graph.

Here it might be an appropriate time to observe that the graph of figure 12.3 is the line-graph of figure 12.3. So we actually have also a Hamiltonian cycle in our original construction.

Chapter 13

Connectivity

To motivate the discussion in this section assume that we have a computer network represented as graph. The vertices are the computer stations and the edges are the links between the computers. Two computers can communicate to each other, if and only if there is a path connecting their vertices in the graph.

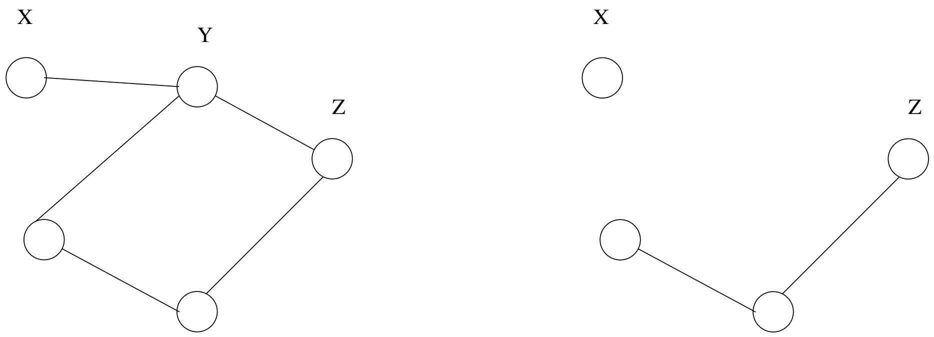

Figure 13.1: Two networks.

The Figures above give two examples of such graphs. Now we can ask yourselves how reliable is this network? Suppose that node Y breaks down in Figure 13.1, then X can not communicate with Z. A more reliable network would be a ring: you have to remove at least 2 nodes to disconnect it.

This problem leads us to a concept in graph theory called vertex connectivity.

Definition 91 *The (vertex) connectivity of a graph G is:*

$$\kappa(G) = min\{|S| : S \subseteq V \quad and \quad G' = (V \setminus S, E) \quad is\ not\ connected\}.$$

A graph $G = (V, E)$ is h(-vertex)-connected if $h \leq \kappa(G)$.

That is, the removal of any $\kappa - 1$ vertices from G does not disconnect G. A subset $S \subseteq V$ that disconnects G is called a *separator*.

That is, our set of vertices can be written in the form:

$$V = A \cup S \cup B$$

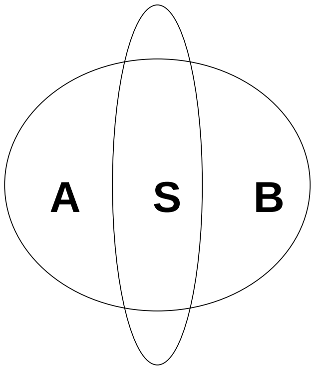

Figure 13.2: An A-B separator S.

with $A \neq \emptyset$ and $B \neq \emptyset$ and no vertex in A is connected to a vertex in B. One problem arising with the definition above is that the complete graph K_n has no separator. For complete graphs we define:

$$\kappa(K_n) = n - 1.$$

Similarly, we can define edge connectivity.

Definition 92 *The edge-connectivity of a graph G is:*

$$\kappa'(G) = min\{|C| : C \subseteq E \quad and \quad G' = (V, E \setminus C) \quad is\ not\ connected\}.$$

A graph $G = (V, E)$ is h-edge-connected if $h \leq \kappa'(G)$.

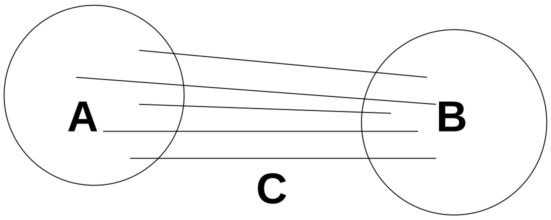

Figure 13.3: An A-B edge separator C.

$C \subseteq E$ is called an *edge separator* if removing the edges in C from the graph will disconnect the graph. Assume that C disconnects the graph into two vertex sets A and B. See Figure 13.3. If C contains edges with both endpoints in A or with both endpoints in B, then removing these edges from C still leaves us an edge separator. Therefore, when interested in the minimum number of edges separating the vertices of the graph into the sets A and B we can consider only the edges with one side in A and the other side in B. We call these edges an A-B cut, and denote them by (A, B).

13.1 Connectivity and degree

Now that we have defined vertex and edge connectivity we state the following inequality.

Theorem 93 *For any graph $G = (V, E)$ we have that $\kappa(G) \leq \kappa'(G) \leq \delta(G)$.*

Proof: To see that that $\kappa'(G) \leq \delta(G)$ let v a vertex of degree δ. Then, clearly $|(\{v\}, V-v)| = \delta$.

To prove the other inequality assume that there is a cut $|(A,B)| = \kappa'$. We show how to choose S to be a set of vertices that contains one endpoint of each edge in (A,B) and $A \not\subseteq S$ and $B \not\subseteq S$. See Figure 13.4.

This is done as follows: if there is a vertex $a \in A$ that is not connected to all vertices in B, then choose the vertices in:

$$S = \{b \mid (a,b) \in (A,B)\} \cup \{u \mid (u,b) \in (A,B) \ \text{ and } \ u \neq a\}.$$

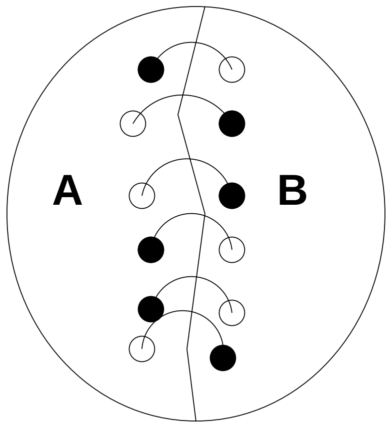

Figure 13.4: The construction of S.

This S is clearly a separator and $\kappa \leq \kappa'$.

If there is no such vertex a, then all nodes in A must be connected to all nodes in B. We get that:

$$\kappa' = |(A,B)| = |A| \cdot |B| \geq |A| + |B| - 1 = n - 1 \geq \kappa$$

By which our proof is complete. $\square$

Example 94 *Let $\epsilon(n,k)$ denote the number of edges in the smallest graph G with n vertices that has connectivity k. We know from the theorem above that $\delta \geq \kappa$. So:*

$$2 \cdot \epsilon(n,k) = \sum_{v \in V} d(v) \geq n \cdot k \quad \Leftrightarrow \quad \epsilon(n,k) \geq \frac{n \cdot k}{2}.$$

This is a good lower bound for $k \geq 2$ as will be shown in one of the exercises in Section 19.

13.2 Menger's Theorem

An important theorem about connectivity is Menger's theorem in its different flavors. Since we are familiar with maximum flows in networks, we will leave the proof of Menger's theorem as an exercise.

Theorem 95 *A graph $G = (V,E)$ is k-edge-connected if and only if there are k edge disjoint paths between every pair of vertices.*

A similar statement can be made about vertex connectivity.

Theorem 96 *A graph $G = (V, E)$ is k-vertex-connected if and only if there are k internally vertex disjoint paths between every pair of vertices. (Internally vertex disjoint means that the internal vertices are disjoint. Clearly, the end vertices cannot be disjoint.)*

To prove these theorems it is easier first to consider directed connectivity. The definitions of directed (edge and vertex) connectivity are analogous to undirected connectivity, only now the edges have directions and a vertex a is connected to b if there is a directed path from a to b. The famous max-flow-min-cut theorem, in the graph with unit edge capacities gives immediatly Menger's theorem (edge version) for directed graphs.

For undirected graphs (edge version) replace each edge by two edges directed in opposite directions. For the vertex version of Menger's theorem, replace each vertex v by two vertices v_{in} and v_{out}, with a single edge from v_{in} to v_{out}. The edges coming into v will be coming into v_{in} and the edges going out from v are connected to v_{out}.

Chapter 14

Matching

A *matching* in a graph $G = (V, E)$ is a subset $M \subseteq E$, such that every two edges in M are independent, or in other words, have different endpoints.

A matching M is said to be *maximal* if there is no larger matching M' such that $M \subseteq M'$. A matching M is a *maximum* matching if there is no larger matching M'. M is called a *perfect* matching if all vertices are matched in M.

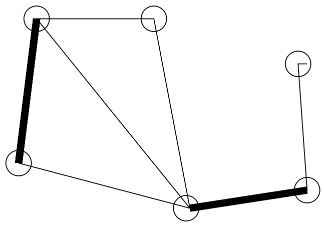

Figure 14.1: The fat edges are a maximal matching, but not a maximum matching.

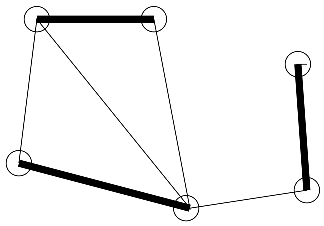

Figure 14.2: The fat edges are a maximum matching, that is also perfect.

14.1 Berge's Theorem

Let M be a matching in $G = (V, E)$. A path P in G is called an M-*alternating* path if its edges alternate between edges of M and edges of $E \setminus M$. An M-alternating path is an M-*augmenting* path if P starts and ends with vertices that are not matched in M.

We have all probably seen that following theorem due to Berge also in some other course.

Theorem 97 (Berge) *M is a maximum matching in a graph $G = (V, E)$ if and only if there are no M-augmenting path in G.*

Proof: If there is an M-augmenting path P, then $(M \cup P) \setminus (M \cap P)$ is a matching of size $|M|+1$. The other way is a bit more difficult. Let M^* be a maximum matching in G and assume that M is a smaller matching. Consider the graph $(V, M \cup M')$. We observe that $d(v) \leq 2$ for all vertices in this graph. Therefore, the graph is a collection of disjoint alternating paths and alternating even cycles. Even cycles and even length path have the same number of edges from both matchings. Since $|M^*| > |M|$, there must be a path that has more edges from M^*. This path is an M-augmenting path. □

14.2 Hall's Theorem

The following theorem due to Hall gives a necessary and sufficient condition to have a matching in a bipartite graph that *saturates* (matches) all vertices on one part.

For a graph G, given $S \subseteq V$, denote by $N(S) = \{u \mid uv \in E \text{ and } v \in S\}$, the set containing all neighbours of S, which are vertices that are adjacent to vertices in S.

Theorem 98 (Hall) *Given a bipartite graph $G = (X \cup Y, E)$, there exists a matching that saturates all vertices in X if and only if for all $S \subseteq X$, $|N(S)| \geq |S|$.*

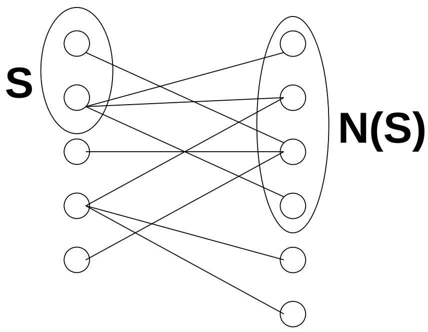

Hall's Theorem

Proof: One direction is trivial: if there is an $S \subseteq X$, such that $|N(S)| < |S|$, then clearly it is not possible to match all vertices in S.

The other direction is a bit trickier. Let M^* be maximum matching and assume that

$$|N(S)| \geq |S|,\ \forall S \subseteq X.$$

Let $x \in X$ be a vertex that is unsaturated in M^* and let Z be the set of all vertices $v \in V$ such that there is an M^*-alternating path from x to v. Let $S = Z \cap X$, $T = Z \cap Y$. The vertices in $S - x$ are matched with the vertices in T and therefore,

$$|S| - 1 = |S - x| = |T|$$

We have also that $N(S) = T$. So:

$$|N(S)| = |T| = |S| - 1$$

in contradiction to our assumption. □

The following corollary is sometimes called the Marriage Theorem.

Corollary 99 *If $G = (X \cup Y, E)$ is a k-regular bipartite graph, $k \geq 1$, then G has a perfect matching.*

Proof: Since G is k-regular, $|X| = |Y|$. Let $S \subseteq X$. Then

$$k \cdot |N(S)| = \sum_{v \in N(S)} d(v) \geq k \cdot |S|,$$

what completes the proof by Hall's Theorem. □

Observe that by removing the perfect matching from a k-regular bipartite graph we get a $(k-1)$-regular bipartite graph. So, by induction, we can partition the edges in the graph into k disjoint perfect matchings.

14.3 König Theorem

Given a graph $G = (V, E)$, we say that $K \subseteq V$ is a *covering* of G if K contains at least one endpoint of each edge of G. K is a *minimum covering* if $|K| \leq |K'|$ for any covering K'.

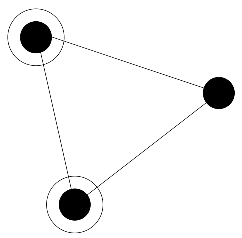

Figure 14.3: A covering of a graph.

We first give some easy properties.

Lemma 100 *If M is a matching and K a covering then $|M| \leq |K|$.*

Proof: Observe that the vertices in K have to cover all the edges in the matching M. □

Lemma 101 *If M is a matching, K a covering and $|M| = |K|$ then M is a maximum matching and K a minimum covering.*

Proof: Let M^* be a maximum matching and K^* a minimum covering. Then we have $|M| \leq |M^*| \leq |K^*| \leq |K|$ and if $|M| = |K|$, then $|M| = |M^*| = |K^*| = |K|$. □

Theorem 102 (König) *Let $G = (V, E)$ be a bipartite graph. Let M^* be a maximal matching and let K^* be a minimal covering. Then $|M^*| = |K^*|$.*

Proof: The proof is similar to the proof of Hall's Theorem. Let $V = X \cup Y$ be the bi-partition of V and let M an arbitrary maximum matching. We will construct a covering K, such that $|K| = |M|$.

Let $U \subseteq X$ be the set of all the vertices in X that are not saturated by the matching M. Let Z be the set of of vertices reachable from U by an M-alternating path and let $S = Z \cap X$, $S = Z \cap Y$. Let $K = T \cup (X \backslash S)$. As in the proof of Hall's Theorem, every vertex in T is saturated by M and $N(S) = T$. We claim that K is a covering for otherwise there is an edge in with one endpoint in S and the other in $Y \setminus T$, contradicting $N(S) = T$. We also have that $|K| = |M|$ and by Lemma 101, the theorem follows. □

14.4 Tutte's Theorem

Now we return to non-bipartite graphs and give a necessary and sufficient condition for a graph to have a perfect matching.

Theorem 103 (Tutte) *A graph $G = (V, E)$ has a perfect matching if and only if $odd(G - S) \leq |S|$ for all $S \subseteq V$, where $odd(H)$ is the number of connected components of H with an odd number of vertices.*

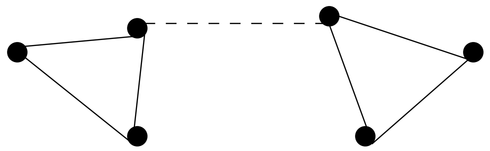

Figure 14.4: $odd(G)$ is 2 without the dashed edge and 0 with it.

Proof: Let M be a perfect matching in G and let S be any subset of V. Let $G_1, G_2, \ldots, G_k$ be the odd components of $G - S$ where $k = odd(G - S)$. We now claim that at least one vertex u_i in $V(G_i)$ is matched in M with some vertex $v_i \in S$. This is clear because the matching is perfect and an odd number of vertices in G_i can not be matched perfectly within G_i and at least one of the vertices from each odd component is matched by an edge to a vertex in S. We then get:

$$odd(G - S) = k = |\{v_1, v_2, \ldots, v_k\}| \leq |S|.$$

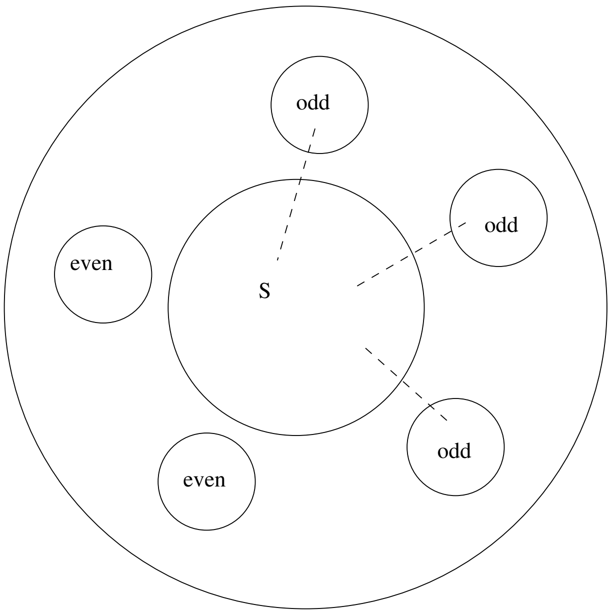

Figure 14.5: The even and odd components in $G - S$.

We prove the other direction by contradiction. Assume that $odd(G - S) \leq |S|$ for all $S \subseteq V$, but G does not have a perfect matching. Let G^* be a maximal graph obtained by adding edges to G as long that G^* does not have a perfect matching. We will obtain the contradiction for G^*. This is as good as showing the contradiction for G since we notice that $odd(G^* - S) \leq odd(G - S) \leq |S|$ for any $S \subseteq V$.

Let U be the set all the vertices in G^* that are connected to all other vertices in G^*. Namely, $d(v) = |V| - 1$ for all $v \in U$. Then U is clearly a complete graph.

We shall later prove that each connected component of $G^* - U$ is a complete graph. Consider the components of $G^* - U$ with even number of vertices. These complete graphs clearly have a perfect matching where all the vertices are matched internally. In components of $G^* - U$ with odd number of vertices, all vertices can be matched internally, but one. Recalling that vertices in U are connected to all vertices of G^*, if there were enough vertices in U, the unmatched vertices from the odd componenets could be matched to vertices in U. But by the assumption, $odd(G^* - U) \leq |U|$, establishing that there is a perfect matching in G^*, since the remaining vertices of U can be matched between themselves, as their number must be even by taking $S = \emptyset$ in the assumption above. See Figure 14.6.

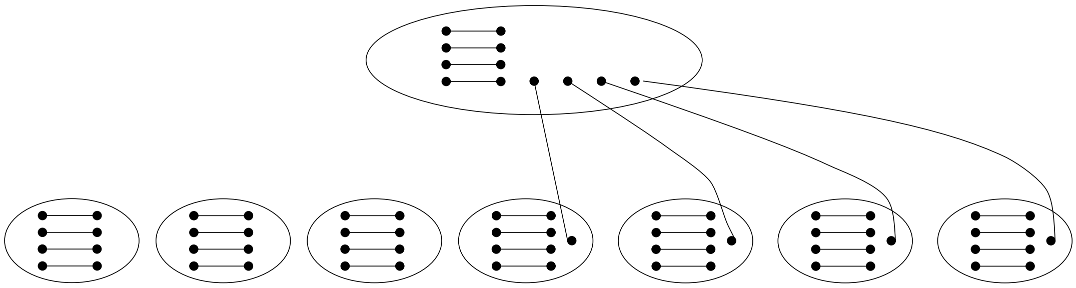

Figure 14.6: The perfect matching in G^*.

It remains only to prove that the connected components of $G^* - U$ are complete graphs. Let H be one of the components of $G^* - U$ which is not complete. We notice that H must contain at least 3 vertices to be a non-complete connected graph. Let H' be a maximal complete subgraph of H and let z be a vertex in H that is not in H', but is connected to some vertex y in H'. Let x be a vertex in H' that is not connected to z. Let $w \in V \setminus U$, be a vertex that is not connected to y. See Figure 14.7.

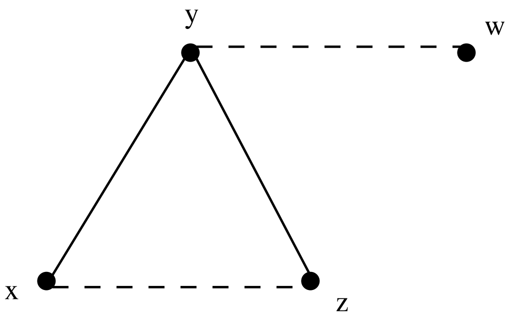

Figure 14.7: The vertices x, y, z, w. Edges are solids and non-edges are dashed.

By the assumption that G^* is maximal, $G^* + xz$ and $G^* + yw$ both have perfect matchings. Call these perfect matchings M_1 and M_2, respectively. Denote by $G' = (V, (M_1 \setminus M_2) \cup (M_2 \setminus M_1))$. Then G' is a collection of alternating even cycles. There are two cases.

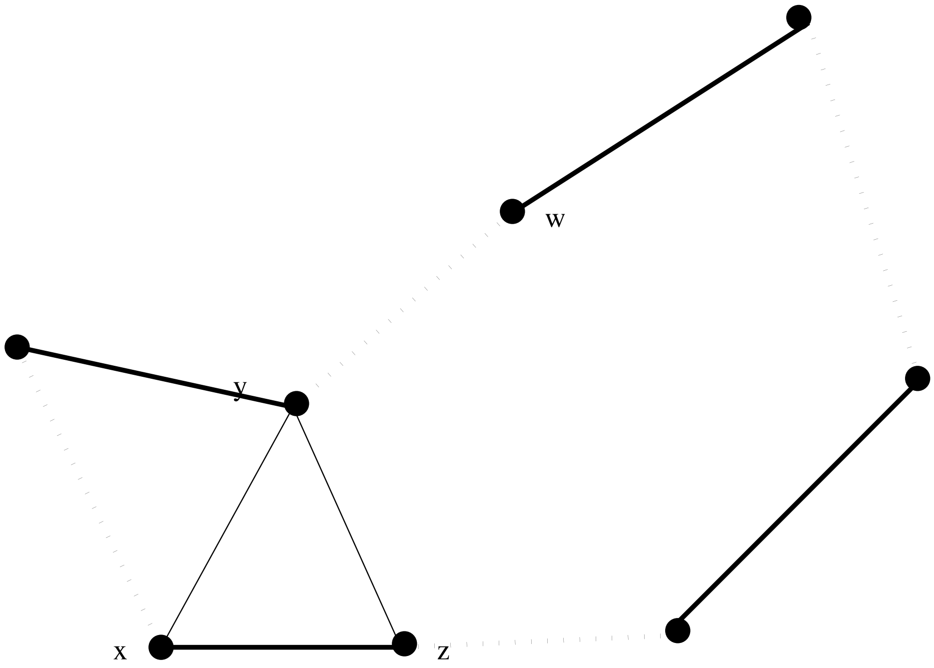

Figure 14.8: xz and yw in the same component of G'. Edges of M_1 are thick and M_2 are dashed.

- xz and yw are in different connected components of G'. Let C be the cycle in G' containing yw. Then the edges of M_1 in C together with the edges of M_2 not in C constitute a perfect matching in G^*.
- xz and yw are in the same connected component of G'. Let C be the cycle in G' containing xz and yw. Assume without loss of generality that the vertices x, y, w, z occur in this order on the cycle C. Then, the edges of M_1 in the section $y, w, \ldots, z$ of C together with the edge yz and the edges of M_2 not in the section $y, w, \ldots, z$ of C constitute a perfect matching in G^*. See Figure 14.8

In both cases we have contradicted the assumption that G^* had no perfect mathing. □

As a corollary to Tutte's theorem we get the following:

Corollary 104 (Petersen's) *If G is a simple 3-regular graph without cut-edges, then G has a perfect matching.*

Proof: From Tutte's theorem we know that G has a perfect matching if $odd(G - S) \leq |S|$ for all $S \subseteq V$. Let S be an arbitrary subset of V, and let $G_1, ..., G_k$ be the odd components of $G - S$. Since G is 3-regular, we know the following two facts:

$$\sum_{v \in V(G_i)} d(v) = 3 \cdot |V(G_i)|$$

and

$$\sum_{v \in S} d(v) = 3 \cdot |S|.$$

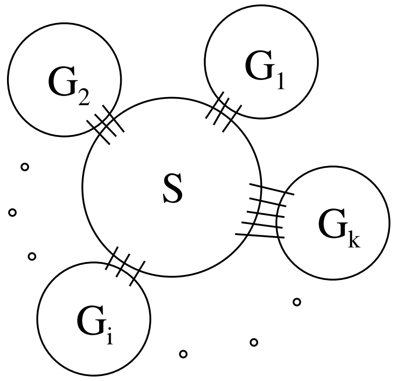

Figure 14.9: The odd components of $G - S$.

where $d(v)$ is the degree of v in the original graph G. Let m_i be the number of edges connecting some vertex in G_i to a vertex in $V - V(G_i)$. Since G_i is not connected to G_j for $i \neq j$, m_i is the number of edges connecting G_i to S. G has no cut-edges, so $m_i \geq 2$. m_i can be expressed in the following way:

$$\begin{aligned} m_i &= \sum_{v \in V(G_i)} d(v) - 2 \cdot |E(G_i)| \\ &= 3 \cdot |V(G_i)| - 2 \cdot |E(G_i)|. \end{aligned}$$

This shows that m_i must be odd since G_i is an odd component. Thus $m_i \geq 3$. Putting all this together we get the wanted:

$$\begin{aligned} odd(G - S) &= k \\ &\leq \frac{1}{3} \sum_{i=1}^{k} m_i \\ &\leq \frac{1}{3} \sum_{v \in S} d(v) \\ &= |S|. \quad \square \end{aligned}$$

Exercise 105 *Consider a graph G with $|V| = 2k$ and $d(v) \geq k + 1$ for all $v \in V$. Then G has a 3-regular spanning subgraph.*

Solution: The degree of every vertex is greater than or equal to half the total number of vertices, so by Dirac's Theorem, G has a Hamiltonian cycle. If we remove every second edge in this cycle, we get a perfect match M, since G has an even number of vertices. The vertices in the new graph $G - M$ each have degree one less than in G. This means, that $G - M$ has a Hamiltonian cycle C, again by Dirac's Theorem. Each vertex in M has degree one, and each vertex in C has degree two. Combining these two we get a 3-regular spanning subgraph $M + C$.

Let $G = (X \cup Y, E)$ be a bipartite graph. The *biadjacency matrix* is a more compact way of representing the adjacency matrix of a bipartite graph. It is defined as an matrix with $|X|$

$$
\begin{array}{c c}
 & \begin{array}{cccc} x_1 & x_2 & \dots & x_n \end{array} \\
\begin{array}{c} y_1 \\ y_2 \\ \vdots \\ y_n \end{array} &
\begin{bmatrix} 0 & 1 & \dots & 0 \\ 1 & 1 & \dots & 1 \\ \vdots & \vdots & & \vdots \\ 0 & 1 & \dots & 0 \end{bmatrix}
\end{array}
$$

Figure 14.10: Example of a biadjacency matrix.

columns and $|Y|$ rows, where entry (i, j) in the matrix is the number of edges between x_i and y_j. See Figure 14.10.

Observe that if a simple bipartite graph G has a perfect matching, then $|X| = |Y|$ and the biadjacency matrix B is a square matrix. Let M be a perfect matching in such a graph. Then, the biadjacency matrix of (V, M) has exactly one 1 is each columns and one 1 in each row, and all other entires of the matrix are zero.

Let $n = |X| = |Y|$. Any permutation $\pi : [n] \to [n]$, corresponds to a perfect matching that matches x_i to $y_{\pi(i)}$, and vice versa. There is such a matching in G, if and only if $B(1, \pi(1)) \cdot B(2, \pi(2)) \cdots B(n, \pi(n)) = 1$, where B is the biadjacency matrix of G. If we wish to count the number of perfect matchings in such a bipartite graph G, it is given by the sum over all permutations π of $[n]$:

$$\sum_{\pi} B(1, \pi(1)) \cdots B(n, \pi(n)).$$

This function of is called the permanent of B and denoted by $per(B)$. The computation of the permanent is a very difficult and important problem (see some complexity course). But the definition above is very similar to the determinant:

$$\det(B) = \sum_{\pi} sgn(\pi) B(1, \pi(1)) \cdots B(n, \pi(n)),$$

where $sgn(\pi) = 1$ if π is a product of an even number of transpositions and $sgn(\pi) = -1$ otherwise. Observe that $per(B) = \det(B) \bmod 2$, and since $\det(B)$ can be computed efficiently, we can find out efficiently if the number of perfect matchings in G is even or odd.

Chapter 15

Edge colouring

Consider the following scheduling problem: in a school, each teacher has to teach a certain number of hours to some classes and each class must be taught a certain number of hours by some teachers. Each teacher can teach only one class at a time and each class can be taught only by one teacher at a time. Ommitting other constraints such as availability of classrooms, etc., this can be modeled nicely as a bipartite graph, with teachers $T_1, ..., T_l$ and classes $C_1, ..., C_p$. An edge (T_i, C_j) indicates that teacher T_i has to teach class C_j for one hour. A matching in this graph corresponds to a possible teaching hour, when some teachers simultaneously teach some classes. The problem is to minimize the number of hours needed for the teachers to teach all their classes?

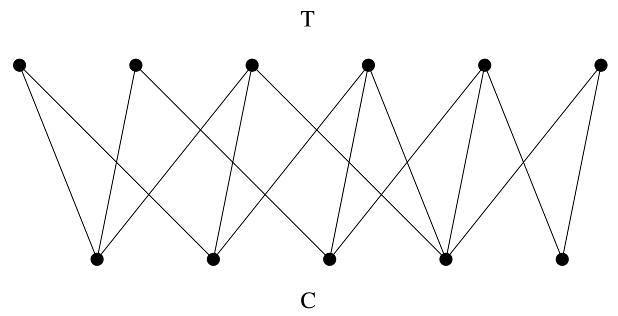

Figure 15.1: A bipartite graph with teachers T and classes C.

Let $G = (V, E)$ be a loopless graph. A *k-edge-colouring* of G is a mapping:

$$C : E \to \{1, ..., k\}.$$

An edge colouring is *proper* if no two edges that are incident to the same vertex are assigned the same colour. The *edge chromatic number* $\chi'(G)$ is defined as the smallest k, such that G has a proper k-edge-colouring. A k-edge-colouring partitions the edges into k disjoint classes in the obvious way:

$$E = E_1 \cup ... \cup E_k, \qquad E_i = C^{-1}(i).$$

Observe that if C is a proper edge colouring of a graph G, then for each colour $i \in [k]$, E_i is a matching in G. It is trivial to see that $\chi'(G) \geq \Delta(G)$, where $\Delta(G)$ is the maximum degree of any vertex in the graph G.

Back to the teacher-class example above, we notice that $\chi'(G)$, for the bipartite teacher-class graph G, is exactly the smallest number of hours required to teach all the hours to the classes. The problem of the smallest number of hours required is completely solved by the following theorem.

Theorem 106 *Let $G = (V, E)$ be a bipartite graph. Then $\chi'(G) = \Delta(G)$.*

Proof: Let $V = X \cup Y$ be the bipartition of the vertices in V. If $|X| \neq |Y|$, then add vertices on the smaller part to have both side of the bipartite graph with equal number of vertices. Assume from now on that $|X| = |Y|$. If the graph G is not Δ-regular, pick one vertex $x \in X$ and another vertex $y \in Y$, with $d(x), d(y) < \Delta$ and add an edge between them. Repeating this step, gives us a Δ-regular bipartite graph.

By the Marriage Theorem, we can find a perfect matching in any regular bipartite graph. Removing the matching from the graph, leaves us with a $(\Delta - 1)$-regular bipartite graph. Repeating this step Δ times, we partition the edges of the graph into Δ matchings, or in other words, we get a proper Δ-edge-colouring. This colouring is clearly a proper colouring also of the edges originally in G. □

15.1 Vizing's Theorem

Vizing Theorem is stated next in its general form. We shall prove it here, however, only for simple graphs.

Theorem 107 (Vizing) $\Delta \leq \chi'(G) \leq \Delta + \mu$, *where μ is the maximum number of parallel edges in G (multiple edges).*

Before we can prove Vizing's Theorem we will need few lemmas.

Lemma 108 *Let G be a connected graph that is not an odd cycle. Then there exists a 2-edge-colouring of G (not necessarily proper) such that both colours are present at all vertices v with $d(v) \geq 2$.*

Proof: The proof is divided into two cases:

G **has an Eulerian circuit.** The Eulerian circuit can be written down in the following way: $v_0 e_1 v_1 \ldots v_0 \ldots e_m v_0$. Define the colouring as follows:

$$C(e_i) = \begin{cases} 1 & \text{if } i \text{ is odd} \\ 2 & \text{if } i \text{ is even} \end{cases}$$

If $d(v) = 2$ for all $v \in V$, then G is an even cycle, and the lemma follows trivially. Otherwise there must exist a vertex v_0 with degree at least 4. The only difficulty here is around v_0, since $C(e_1)$ might be equal to $C(e_m)$. But this is not a problem, since v_0 also appears somewhere in the middle of the circuit.

G **has no Eulerian circuit.** Add a new vertex u to the graph. Connect u to all $v \in V$ with $d(v)$ odd. There must be an even number of these vertices making the degree of u even. Since all the vertices of this new graph have even degree, it has an Eulerian circuit and can be coloured as above. This colouring is also a 2-edge-colouring of G. We claim that both colours are incident to all vertices v with $d(v) \geq 2$. To see this observe that if $d(v)$ was even, then this is clearly so. If $d(v)$ was odd and $d(v) > 2$, then apart from the edge added between v and u, there are two other edges incident to v that appear consecutively on the Euler circuit. These edges were coloured with the 2 colours. □

Before we prove Vizing's theorem we need the following notation. Let C be an edge colouring of G. Define $c(v)$ to be the number of different colours present at vertex v. Clearly $c(v) \leq d(v)$. C is proper if and only if $c(v) = d(v)$ for all $v \in V$. Let C and C' be k-edge-colourings of G. C' is said to *improve* C if

$$\sum_{v \in V} c'(v) > \sum_{v \in V} c(v).$$

C is an *optimal k-edge-colouring* if no C' improves it.

Lemma 109 *Let $G = (V, E)$ be a simple graph and let $C = (E_1, E_2, \cdots, E_k)$ be an optimal k-edge-colouring of G. Let v be a vertex in G such that there exist two colours i and j fulfilling that E_i has two edges incident to v and no edge in E_j is incident to v. Then the connected component in $(V, E_i \cup E_j)$ containing v is an odd cycle.*

Proof: Let H be the connected component containing v in the subgraph $(V, E_i \cup E_j)$ of G. Suppose that H is not an odd cycle. By Lemma 108, we can colour H with two colours i and j, such that for any vertex with degree 2 or more in H, both i and j are present at the vertex. Let $C' = (E'_1, E'_2, \cdots, E'_k)$ be a new colouring of G that is equal to the old colouring C except from the edges in the connected component of v in H that are coloured according to new colouring. This colouring C' must fulfill $c'(v) = c(v) + 1$ and $\forall u \in V$: $c'(u) \geq c(u)$, contradicting the assumption that C was an optimal k-edge-colouring. □

We give next a different proof of Theorem 106 using the notation we have just developed.

Theorem 110 *Let $G = (V, E)$ be a bipartite graph. Then $\chi'(G) = \Delta$.*

Proof: Suppose that $\chi' > \Delta$ and let C be an optimal Δ-edge-colouring. Since C can not be a proper colouring, there exists a vertex v such that $c(v) < d(v) \leq \Delta$. For this vertex, at least two incident edges have the same colour i, and there is some colour j not present at v. By Lemma 109, G must have an odd cycle, contradicting that it is a bipartite graph. □

We are now ready to prove the abridged version of Vizing's Theorem.

Theorem 111 (Vizing) *If G is a simple graph, then $\chi'(G) = \Delta$ or $\chi'(G) = \Delta + 1$*

Proof: By the definition clearly $\chi' \geq \Delta$. Suppose that $\chi' > \Delta + 1$. Let C be an optimal edge-colouring with $\Delta + 1$ colours. As in the proof of Theorem 110 there exists a vertex v with $c(v) < d(v)$ and again, at least two edges incident to v are coloured with the same colour i_1. Let w and u_1 be the other endpoints of these two edges. Pick a colour i_2 not present at u_1. If this colour is present at v, we let u_2 denote the other endpoint of the edge coloured i_2. Repeat picking a colour i_{l+1} not present at u_l and if present at v, name the other new endpoint u_{l+1}. See Figure 15.2.

There are two cases stopping this process. The first case is that a colour i_{k+1} which is not present at u_k is also not present at v. In this case we change the colouring C by shifting

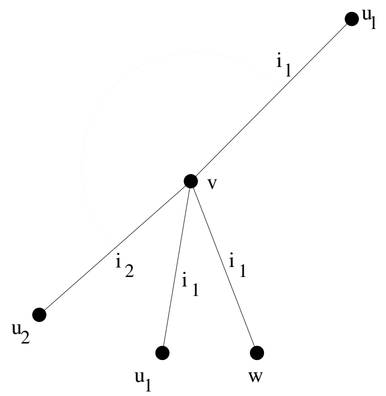

Figure 15.2: Situation in proof of Vizing's theorem.

the colours of the edges $u_i v, 1 \leq i \leq l$, such that $u_i v$ is coloured with i_{i+1}. See Figure 15.3. The new colouring C' must fulfill that $c(v) < c'(v)$ since the colour i_{l+1} has been added at v. $c(u) \leq c'(u)$ for all other vertices u in G, since the recolouring only affects the vertices u_i, which we carefully recoloured with a colour not present, preserving the inequality. But then C' improves C contradicting that C is optimal.

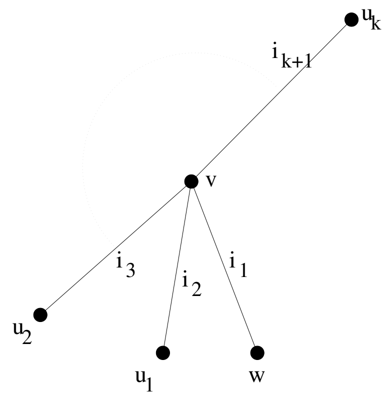

Figure 15.3: Recolouring in the first case

The other case is that the colour $i_{l+1} = i_k$ for some $k < l + 1$. We construct two new optimal edge-colourings C' and C''. The colouring C' is obtained from C as before by shifting the colours of the edges $u_1 v, \ldots, u_{k-1} v$ such that they have the colours $i_2 \cdots i_k$ respectively. As argued in the first case, the number of colours $c(u_i)$ and $c(v)$ present at u_i and at v are not decreased, ensuring that C' is indeed an optimal colouring. C'' is obtained by colouring the edges $u_1 v, \ldots, u_{l-1} v$ with the colours $i_2 \cdots i_l$ and edge $u_l v$ coloured i_k. Let i_0 be a colour not present at v. Such colour always exists since $d(v) \leq \Delta$ and we have $\Delta+1$ colours. Consider the two subgraphs $(V, E'_{i_0} \cup E'_{i_k})$ and $(V, E''_{i_0} \cup E''_{i_k})$ (E' and E'' are the edge-partitions according to the optimal colourings C' and C'', respectively). The components containing v in each of these subgraphs are odd-cycles by Lemma 109, or otherwise the colourings C, C' and C'' are not optimal. See Figures 15.4 and 15.5. But since the degree of u_k must be 2 in the first subgraph and only 1 in the second subgraph, this is impossible. $\square$

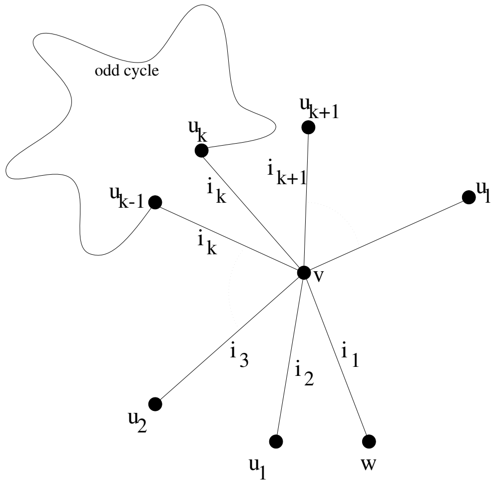

Figure 15.4: The odd cycle of C'

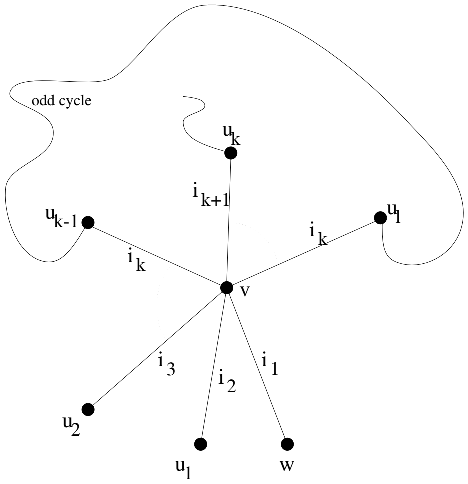

Figure 15.5: The odd cycle of C''

Chapter 16

Cliques

Let $G = (V, E)$ be a simple graph. A subset $S \subseteq V$ is a *clique* if $\forall u, v \in S : uv \in E$. A subset $S \subseteq V$ is an *independent set* if $\forall u, v \in S : uv \notin E$. Recall that a *(vertex-) covering* of G is a set $S \subseteq V$, such that every edge in E is incident to some vertex in S.

Theorem 112 *S is an independent set in $G = (V, E)$ if and only if $V \setminus S$ is a covering of G.*

Proof: Let S be an independent set and suppose that $S^c = V \setminus S$ is not a covering. Let uv be an edge not covered by S^c. It follows that both u and v are in S, contradicting that S is independent.

For the other direction let S^c be a covering and suppose that $S = V \setminus S^c$ is not an independent set. Let uv be an edge connecting two vertices u and v in S. This edge must be covered by either u or v in S^c, contradicting that $S = V \setminus S^c$. □

S is a *maximum independent set* if for any independent set S': $|S| \geq |S'|$. Similarly a *minimum covering* is a covering K such that for any other covering K': $|K| \leq |K'|$. Let $\alpha(G)$ denotes the size of the maximum independent set in G and $\beta(G)$ is the size of the minimum covering of G.

Corollary 113 $\alpha + \beta = |V|$.

Proof: Let K be a minimal covering of G, that is $|K| = \beta$. We have that $V \setminus K$ is independent and, thus $|V \setminus K| \leq \alpha$, implying that $\alpha + \beta \geq |V|$. The other direction is treated similary. □

The analogies of vertex-independence and vertex-covering for edges are *matching* and *edge-covering*, respectively. In previous sections the notion of matchings has already been introduced. An *edge-covering* is a subset $C \subseteq E$ for which every vertex v in V is incident to an edge in C. We denote by $\alpha'(G)$ and $\beta'(G)$ the sizes of the maximum matching and minimum edge-covering, respectively. It is clear that β' is well-defined only for a graph with $\delta \geq 1$. We therefore restrict ourselves to such graphs in the discussion about edge-covering.

Figure 16.1 shows that an analogy of Theorem 112 for matchings and edge-coverings does not hold. However, we can prove the following theorem.

Theorem 114 $\alpha' + \beta' = |V|$

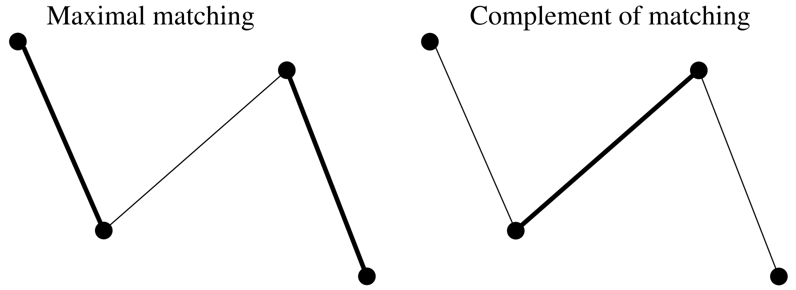

Figure 16.1: The complement of a matching is not an edge-covering.

Proof: Let M be a maximum matching and U be the set of M-unsaturated vertices. Then $|M| = \alpha'$. Let $E' \subseteq E$ be an edge-cover of U, such that $|E'| = |U|$. We have

$$\beta' \leq |M \cup E'| = |M| + |E'| = |M| + |U| = \alpha' + (|V| - 2\alpha')$$

and thus $\alpha' + \beta' \leq |V|$.

Let L be a minimum edge cover. Then, $|L| = \beta'$. Let M be a maximum matching in the subgraph (V, L) and let U be the M-unsaturated vertices. We have

$$|L| - |M| = |L \setminus M| \geq |U| = |V| - 2|M|,$$

and we obtain $\beta' + \alpha' \geq |L| + |M| \geq |V|$. $\Box$

From the theorems above we obtain the following equation: $\alpha + \beta = \alpha' + \beta' = |V|$. By Königs theorem we know that in a bipartite graph $\alpha' = \beta$, from which the following theorem follows.

Theorem 115 *In a bipartite graph, the size of a maximum independent set equals the size of the minimum edge-covering. That is $\alpha = \beta'$.*

16.1 Ramsey's Theorem

A question raised by Ramsey, which so-to-speak founded Ramsey's Theory, was: does there exist an integer $r(k, l)$ such that any graph with $r(k, l)$ or more nodes has a k-clique *or* an l-independent set (or both). The answer is yes, and we will get some bounds on the values of $r(k, l)$ in this section. We start out with a few examples.

Example 116 *Clearly $r(1, l) = 1$ and $r(k, 1) = 1$. Furthermore $r(2, l) = l$ (if the graph does not contain an l-independent set, then it definitely has an edge—i.e. a clique of size 2), and $r(k, 2) = k$ (similarly). We may thus fill-in the first entries in the corresponding table:*

		l					
$r(k,l)$		1	2	3	4	5	$\cdots$
k	1	1	1	1	1	1	$\cdots$
	2	1	2	3	4	5	$\cdots$
	3	1	3	?	?		
	4	1	4	?			
	5	1	5				
	$\vdots$	$\vdots$	$\vdots$				

To see that $r(3,3) \geq 6$, we can give an example of a graph on 5 vertices that does not contain a 3-cycle nor a 3-independent set. Such a graph is the cycle on 5 vertices. We will show soon that in fact $r(3,3) = 6$.

Theorem 117 *$r(k,l)$ is well defined and $r(k,l) \leq r(k-1,l) + r(k,l-1)$.*

Proof: By induction on k and l. For k or $l \leq 2$, the inequality holds by the above example. Let $G = (V, E)$ be a graph on $r(k-1,l) + r(k,l-1)$ vertices. Consider an arbitrary vertex v of G and its neighbours $N(v) := \{u \mid \{u,v\} \in E\}$.

If $|N(v)| \geq r(k-1,l)$, then the induced graph $G[N(v)]$ has by the inductive hypothesis an l-independent set or a $(k-1)$-clique. If it has an l-independent set, then it is also an independent set in G and we are done. If it has a $(k-1)$-clique, this clique together with v forms a k-clique in G as required.

On the other hand, if $|N(v)| < r(k-1,l)$ then $|V \setminus N(v) \setminus \{v\}| \geq r(k,l-1)$ and the same arguments apply switching the roles of neighbours and non-neighbours of v. $\Box$

Let us dwell one more moment at this and show that if both $r(k-1,l)$ and $r(k,l-1)$ are even, then $r(k,l) \leq r(k-1,l) + r(k,l-1) - 1$. By Corollary 56, we know that in every graph there will be an even number of vertices with odd degree. Letting G be a graph with $r(k-1,l) + r(k,l-1) - 1$ vertices, we can pick a vertex v with even degree. Now, $|N(v)|$ is even and thus $|N(v)| \geq r(k-1,l)$ or $|N(v)| \leq r(k,l-1) - 2$. The latter will again imply $|V \setminus N(V) \setminus \{v\}| \geq r(k,l-1)$.

Corollary 118 $r(k,l) \leq \binom{k+l-2}{k-1}$.

Proof: By induction on the pair (k,l). If $k, l \leq 2$ then this is easy to verify. Otherwise, we get

$$r(k,l) \leq r(k-1,l) + r(k,l-1) \leq \binom{k+l-3}{k-2} + \binom{k+l-3}{k-1} = \binom{k+l-2}{k-1}. \quad \Box$$

As a consequence we observe that $r(k,k) \leq 2^{2k-2}$. We can also add the following lower bound on $r(k,k)$:

Theorem 119 (Erdös) $r(k,k) \geq 2^{k/2}$.

Proof: The claim can be easily verified for $k = 1, 2$. So assume that $k \geq 3$. We give a non-constructive proof by the so called *probabilistic method.* Recall that there are $2^{\binom{n}{2}}$ simple labeled graphs with n vertices, and let G be a randomly chosen graph. We shall establish the following fact: the probability of G having a clique of size k *or* an independent set of size k is strictly smaller than 1, when $n < 2^{k/2}$; written

$$\Pr(G \text{ has } K_k \textit{ or } G \text{ has } K_k^c) < 1.$$

Then, for any $n < 2^{k/2}$, there must be some graph G on n vertices that has no clique and no independent set of size k, establishing that $r(k,k) \geq 2^{k/2}$.

Consider the probability of the event that G has a clique of size k.

$$\begin{aligned}
\Pr(G \text{ has } K_k) &\leq \frac{\binom{n}{k} 2^{\binom{n}{2}-\binom{k}{2}}}{2^{\binom{n}{2}}} \\
&= \binom{n}{k} 2^{-\binom{k}{2}} \\
&\leq \frac{n^k 2^{-\binom{k}{2}}}{k!} \\
&< \frac{2^{k^2/2} 2^{-\binom{k}{2}}}{k!} \qquad \text{for } n < 2^{k/2} \\
&= \frac{2^{k/2}}{k!} \\
&< \frac{1}{2} \qquad \text{for } k \geq 3
\end{aligned}$$

And since cliques and independent sets are symmetric, similar arguments apply for independent sets. Together this gives us

$$\Pr(G \text{ has } K_k \textit{ or } G \text{ has } K_k^c) \leq \Pr(G \text{ has } K_k) + \Pr(G \text{ has } K_k^c) < 1,$$

as required. □

This is a non-constructive results; the best constructive lower bound that are known are much smaller. See also Part III.

The precise value of $r(k, l)$ is known only for very few values of k and l. In Figure 16.2 we give four other values of $r(k, l)$ that are known precisely. The upper bounds in these cases can be derived from Theorem 117 and the remark following it.

		l					
$r(k,l)$		1	2	3	4	5	$\cdots$
k	1	1	1	1	1	1	$\cdots$
	2	1	2	3	4	5	$\cdots$
	3	1	3	$\boxed{6}$	$\boxed{9}$		
	4	1	4	$\boxed{9}$	$\boxed{18}$		
	5	1	5				
	$\vdots$	$\vdots$	$\vdots$				

Figure 16.2: r(k,l).

16.2 Turán's Theorem

In this section we consider the following question: suppose that we know that a given simple graph G has no K_{h+1}-subgraph; what it the maximum number of edges that G might have? For $h = 2$ one can easily see that a bipartite graph $G = (V, E)$ with the bipartition $V = X \cup Y$

and the edges $E = X \times Y$ does not contain K_3 (it contains no odd cycles of any length). A simple calculation shows that among all such bipartite graphs, the one that has the largest number of edges is the one in which the vertices of V are divided equally between X and Y.

We can generalise this: A graph $G = (V, E)$ is said to be *h-partite* if its vertices are partition into disjoint independent sets $V = V_1 \cup V_2 \cup \ldots \cup V_h$ and $V_i \cap V_j = \emptyset$, $i \neq j$. A *complete h-partite* graph is an h-partite graph $G = (V_1 \cup V_2 \cup \ldots \cup V_h, E)$ where $E \supseteq V_i \times V_j$, $i < j$. Thus, in a complete h-partite graph G, $|E| = \sum_{i<j} |V_i| \cdot |V_j|$.

The following fact follows immediately from the definitions above.

Fact 120 *An h-partite graph has no $(h+1)$-clique.*

Given two graphs G and H, the *join* of G and H, written $G \vee H$, is defined as

$$G \vee H := (\, V(G) \cup V(H),\; E(G) \cup E(H) \cup V(G) \times V(H)\,).$$

Observe that the complete h-partite graph with $V = V_1 \cup V_2 \cup \ldots \cup V_h$ is precisely $G = G_1 \vee G_2 \vee \ldots \vee G_h$, where $G_i = (V_i, \emptyset)$.

Lemma 121 *Let $T_{n,h}$ be the complete h-partite graph with n vertices and the h-parts being of almost equal sizes. Namely, $V(T_{n,h}) = U_1 \cup U_2 \cup \ldots \cup U_h$, $\lfloor \frac{n}{h} \rfloor \leq |U_i| \leq \lceil \frac{n}{h} \rceil$. Then, for any given simple h-partite graph G with n vertices, $|E(G)| \leq |E(T_{n,h})|$.*

Proof: Let $G = (V_1 \cup V_2 \cup \ldots \cup V_h, E)$ be an h-partite graph. Assume without loss of generality that G is complete and that $|V_1| \leq |V_2| \leq \cdots \leq |V_h|$. Observe that if $|V_h| \leq |V_1| + 1$ then $G \cong T_{n,h}$ and there is nothing to prove. So suppose that $|V_h| \geq |V_1| + 2$. Clearly $|E| \leq \sum_{i<j} |V_i| \cdot |V_j|$. Take any vertex $x \in V_h$ and set $V_1' = V_1 \cup \{x\}$, $V_h' = V_h \setminus \{x\}$ and $V_i' = V_i$, $i = 2, 3, \ldots, h-1$.

Then

$$\begin{aligned}
\sum_{i<j} |V_i'| \cdot |V_j'| &= (|V_1| + 1)(|V_2| + |V_3| + \cdots + |V_{h-1}|) \\
&\quad + (|V_1| + 1)(|V_h| - 1) \\
&\quad + \sum_{1<i<j<h} |V_i| \cdot |V_j| \\
&\quad + (|V_h| - 1)(|V_2| + |V_3| + \cdots + |V_{h-1}|) \\
&= \sum_{i<j} |V_i| \cdot |V_j| + |V_h| - |V_1| - 1 \\
&= |E| + |V_h| - |V_1| - 1 \\
&\geq |E| + 1 \\
&> |E|.
\end{aligned}$$

The same process can be repeated until we get $T_{n,h}$. $\square$

Observe that in $T_{n,h}$, $|V_i| = \lfloor \frac{n+i-1}{h} \rfloor$, $i = 1, \ldots, h$.

Suppose that G and H are two graphs on n vertices and let $d_1 \leq \cdots \leq d_n$ and $f_1 \leq \ldots \leq f_n$ be the degree sequences of their vertices, respectively.. We say that H *degree-majorises* G if $f_i \geq d_i$, for $i = 1, \ldots, n$.

We only need one last lemma for the proof of Turán's Theorem.

Lemma 122 *Let $G = (V, E)$ be a simple graph and assume that G contain no K_{h+1} as a subgraph. Then there exists a h-partite graph H that degree majorizes G.*

Proof: The proof proceeds by induction on h. For $h = 1$, the assumption that G contains no K_2 means that G has no edges, so it is 1-partite (an independent set).

Assume that the claim holds for all graphs that contain no K_h as subgraphs and assume that G contains no K_{h+1}. Let $v \in V$ be a vertex such that $d(v) = \Delta(G)$. Let $V_1 = N(v)$ (so $|V_1| = \Delta(G)$) and $G_1 = G[V_1]$ the subgraph of G induced by V_1.

Since G has no clique of size $h+1$, G_1 has no clique of size h, as by construction this would give a clique of size $h+1$ in G. So by the induction hypothesis there exists a $(h-1)$-partite graph H_1 the degree majorizes G_1. Now let $G_2 = (V \setminus V_1, \emptyset)$. Clearly, for $u \in V_1$ we have that $N_G(u) \subseteq N_{G_1 \vee G_2}(u)$, and for $w \in V \setminus V_1$ we have that $d_{G_1 \vee G_2}(w) = \Delta(G) \geq d_G(w)$, so $G_1 \vee G_2$ degree majorizes G.

Finally, we get that $H_1 \vee G_2$ degree majorizes $G_1 \vee G_2$, and since by induction hypothesis H_1 is $(h-1)$-partite, we got an h-partite graph $H_1 \vee G_2$ that degree majorizes G. □

We are now ready to prove Turán's Theorem.

Theorem 123 (Turán) *Let $G = (V, E)$ be a simple graph and assume that G contains no K_{h+1} as a subgraph. Then,*

$$|E| \leq \sum_{0 \leq i < j < h} \lfloor \frac{|V| + i}{h} \rfloor \lfloor \frac{|V| + j}{h} \rfloor$$

Proof: By Lemma 122, G is degree-majorized by an h-partite graph H and by Lemma 121 we get that:

$$|E(G)| \leq |E(H)| \leq |E(T_{|V(G)|,h})| \leq \sum_{0 \leq i < j < h} \lfloor \frac{V(G) + i}{h} \rfloor \lfloor \frac{V(G) + j}{h} \rfloor. \quad \square$$

Chapter 17

Vertex colouring

Let $G(V, E)$ be an undirected graph. A map $C : V \mapsto [k]$ is called a *k-colouring* of G (or a *k-vertex-colouring*). If C further satisfies the condition: $uv \in E \Rightarrow C(u) \neq C(v)$, then it is called a *proper k-colouring.* We say that G is *k-colourable* if it has a proper k-colouring. Finally, we define the *chromatic number* of G, $\chi(G) = min\{k \mid G \text{ is } k\text{-colourable}\}$.

Unless otherwise stated, we shall always mean *proper colouring* when we talk about colouring. Clearly, if G contains self-loops then it has no proper colouring. Also, multiple edges do not make any difference for colouring. Therefore, we shall only deal with simple graphs in this section.

Remark: Suppose that C is a colouring of a graph G. If we look at the set of vertices of same colour, i.e. $C^{-1}(i)$, we see that it is an independent set. So, a colouring is a partition of the vertices V into independent sets and a graph is k-colourable if and only if it is k-partite.

Theorem 124 *Let G be a simple graph. Then, $\chi(G) \leq \Delta(G) + 1$.*

Proof: By induction on the size of V. For $|V| = 1$, G is trivially 1-colourable. Let $v \in V$, such that $d(v) = \Delta$ and let $G' = G - v$. By the inductive hypothesis G' is $(\Delta(G') + 1)$-colourable. But $\Delta(G') \leq \Delta(G)$ and therefore G' is $(\Delta + 1)$-colourable. Since v only has Δ neighbours, we have at least one colour not colouring any of v's neighbours, which we can use for v to get a proper $(\Delta + 1)$-colouring of G. □

It is easy to see that the complete graph K_n and an odd cycle are examples where we need $\Delta + 1$ colours. The following theorem that is stated without a proof claims that these are the only such examples.

Theorem 125 (Brooks) *If G is a simple graph that is not a complete graph and not an odd cycle, then $\chi(G) \leq \Delta$.*

17.1 Chromatic polynomials

Given a graph G and an integer $k \geq 0$, define the function $\Pi_G(k)$ to be the number of different proper k-colourings of G. By definition, we immediately get that if $0 \leq k < \chi(G)$, then $\Pi_G(k) = 0$.

Example 126 *For the independent set K_n^c we get: $\Pi_{K_n^c}(k) = k^n$, since we have no restrictions on what colours each vertex can be given.*

Example 127 *For the complete graph K_n we get: $\Pi_{K_n}(k) = k(k-1)\cdots(k-n+1)$. Notice that if $k < n$, then $\Pi_{K_n}(k) = 0$.*

Exercise 128 *What is $\Pi_G(k)$ if G is a tree? If G is a cycle?*

The functions $\Pi_G(k)$ are called *chromatic polynomials* since as we will show next, they are indeed polynomials in k. Recall the notion of edge-contraction that was defined in Section 11 and is used here in a similar manner.

Theorem 129 *Let $G = (V, E)$ be graph and $k \geq 0$. Then for any $e \in E$,*

$$\Pi_G(k) = \Pi_{G-e}(k) - \Pi_{G \cdot e}(k).$$

Proof: Let $e = (u, v)$. The different clourings of $G - e$ can be partitioned into two types: colourings with u and v having different colours and colourings with u and v having the same colour. Only the first type of colourings are also proper colourings of G. But colourings of the second type where u and v have the same colours correspond precisely to colourings of $G \cdot e$. Therefore,

$$\Pi_G(k) = \Pi_{G-e}(k) - \Pi_{G \cdot e}(k),$$

completing the proof. □

The following corollary justifies the name *chromatic polynomials* and can be proved by induction. The proof is left as an exercise. The recurrence established in the last theorem can be used to compute $\Pi_G(k)$, where Example 126 is used if G has no more edges. Computing the chromatic polynomial of a graph G is not an easy task, since it give us the chromatic number of G, which is an NP–complete problem.

Corollary 130 *$\Pi_G(k)$ is a polynomial in k of degree $|V(G)|$, with integer coefficients.*

The chromatic polynomials $\Pi_G(k)$ have been defined for $k \geq 0$. Surpringly, the value of these polynomials at $k = -1$ has an interesting meaning. Recall that we used $aco(G)$ to denote the number of acyclic orientations of an undirected graph G.

Theorem 131 *Let $G = (V, E)$ be an undirected graph. Then,*

$$aco(G) = (-1)^{|V|} \cdot \Pi_G(-1).$$

Proof: We will first show that $aco(G) = aco(G - e) + aco(G \cdot e)$, for any edge $e \in E$. Let $e = (u, v)$ and consider an acyclic orientation of $G - e$. If there is a directed path from u to v, or from v to u, in the orientation of $G - e$, then e can be directed only in one direction to get an acyclic orientation of G. If there is no such path, then e can be oriented in both direction to get two different acyclic orientations of G. But the orientation of $G - e$ for which e can be directed in either direction correspond precisely to acyclic orientations of $G \cdot e$, establishing that $aco(G) = aco(G - e) + aco(G \cdot e)$. We can now prove the claim by induction on the number of edges in G. By Example 126, for graphs G with no edges, we have:

$$aco(G) = (-1)^{|V|} \cdot \Pi_G(-1) = (-1)^{2 \cdot |V|} = 1.$$

Assuming the the claim holds for any graph with $|E|-1$ edges, we get by Theorem 129:

$$\begin{aligned} aco(G) &= aco(G-e) + aco(G\cdot e) \\ &= (-1)^{|V(G)|}\cdot \Pi_{G-e}(-1) + (-1)^{|V(G)|-1}\cdot \Pi_{G\cdot e}(-1) \\ &= (-1)^{|V(G)|}\cdot (\Pi_{G-e}(-1) - \Pi_{G\cdot e}(-1)) \\ &= (-1)^{|V(G)|}\cdot \Pi_{G}(-1), \end{aligned}$$

establishing the claim. □

17.2 Girth

If asked to construct graphs with arbitrarily large chromatic numbers one might take the complete graphs K_n. These graphs, however, have many edges and one might consider the problem of constructing graphs with large chromatic numbers with fewer edges. From a first glance, it is not clear at all if there are graphs with arbitrarily large chromatic numbers that do not have very short cycles. It is not clear even if there graphs with arbitrarily large chromatic numbers and no triangles. We shall try to asnwer some of these questions now. The *girth* of a graph G is defined to be the length of the shortest cycle in G.

The following theorem is stated without a proof.

Theorem 132 *Given any positive integers l and k, there exist a graph with girth $\geq l$ and $\chi \geq k$.*

We prove here only the following special case for girth $= 4$.

Theorem 133 *For any $k \geq 1$ there is a graph G_k with $\chi(G_k) = k$ and no cycle of length* 3.

Proof: The proof goes by induction in the number of colours k. For $k = 1, 2, 3$, it is easy to construct G_1, G_2 and G_3 being a graph with one vertex and no edges, a graph with two vertices and a single edge connecting them, and a graph with five vertices in a cycle, respectively.

Assume that we have a graph $G_k = (V_k, E_k)$ with no cycles of length 3 and $\chi(G_k) = k$, and construct $G_{k+1} = (V_{k+1}, E_{k+1})$ with the same properties. Denote $V_k = \{v_1, \ldots, v_n\}$ and let

$$V_{k+1} = \{v_1, v_2, \cdots, v_n, u_1, u_2, \cdots, u_n, w\}.$$

The edge set E_{k+1} is defined as follows. See Figure 17.1.

- The vertices v_i are connected as they were in G_k.
- Each vertex u_i is connected to all $N(v_i)$.
- w is connected to each vertex u_i.

Now we have two things to show: first that there are no cycle of length 3 in G_{k+1} and second that $\chi(G_{k+1}) = k+1$. By the inductive hypothesis, there are no cycles of length 3 in G_k. So there can be at most two v_i's in a 3-cycle. To see that there can only be one u_i in a 3-cycle, and that w cannot be in 3-cycle, observe that there are no edges between any

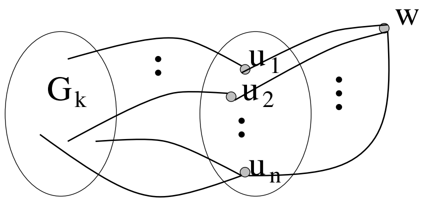

Figure 17.1: The construction of G_{k+1}.

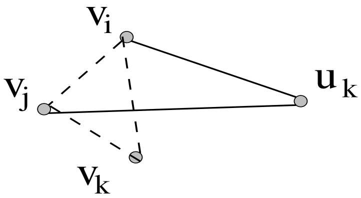

Figure 17.2: Cycles in G_{k+1}.

pair of u_i's and that w is not adjacent to any v_i. Therefore, any 3 cycle in G_{k+1} must be of the form $v_i v_j u_k v_i$. But then, $v_i v_j v_k v_i$ would also be a cycle in G_{k+1} since u_k is connected to all the neighbours of v_k, in contradiction to the assumption. Therefore, there can not be any 3-cycle in G_{k+1}. See Figure 17.2.

It remains to prove that $\chi(G_{k+1}) = k+1$. Clearly $k = \chi(G_k) \leq \chi(G_{k+1}) \leq k+1$. Assume by way of contradiction that G_{k+1} is k-colourable and colour it with k colours. Renumber the colours so that w is coloured with colour k. Then, $u_1, \ldots, u_n$ are colours with colours from $[k-1]$. If a vertex v_i has the colour k, then change that colour to the colour of u_i. This colouring is also proper as none of the vertices $N(u_i)$ and thereby none of the vertices $N(v_i)$ has the same colour as u_i. But now we have reached a contradiction to the fact that $\chi(G_k) = k$, since the colouring of G_{k+1} induces a proper colouring of the subgraph G_k with $k-1$ colours. □

Chapter 18

Planar graphs

A graph is *planar* if it can be *drawn*, or in other words *embedded*, in the plane so that its edges do not intersect each other, except in their ends. See Figure 18.1. A specific embedding of a planar graph is called a *plane* graph. Notice that a planar graph may have several different plane graphs that are isomorphic to it.

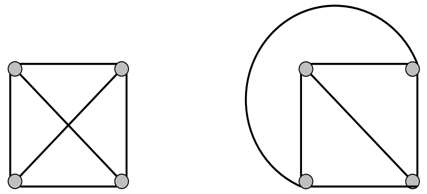

Figure 18.1: A non-planar a planar drawings of K_4.

A *Jordan curve* in the plane is a closed continuous curve that does not intersect itself. See Figure 18. To be able to reason about planar graphs we will need the following theorem about the topology of the plane. The theorem should be intuitively clear and we shall not prove it here. (The proof, however, is far from being trivial.)

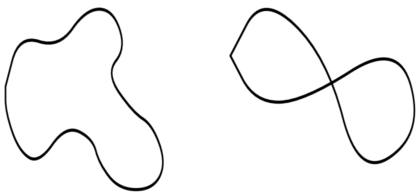

Figure 18.2: The curve on the left is a Jordan curve. The curve on the right intersect itself and therefore it is not a Jordan curve.

Theorem 134 (Jordan) *A Jordan curve J partitions the plane into two parts: the* inside *and the* outside *(or* interior *and* exterior*). The intersection of the two parts is exactly the jordan curve J. Any continuous line connecting a point inside J to a point outside J must intersect J.*

Before we give an example of how to use the theorem in connection with planar graphs we state the following fact.

Fact 135 *Cycles in a graph G are Jordan curves in any planar embedding of G.*

Lemma 136 *K_5 is non planar.*

Proof: Look at Figure 18.3. We have named the vertices $v_1, \ldots, v_5$. Consider the cycle $v_1v_2v_3v_1$. Then, v_4 can be inside or outside the Jordan curve formed by this cycle. As the two cases are symmetric, assume that v_4 is inside the triangle (the cycle). The remaining vertex v_5 can be in one of the three small triangles (named a, b and c in the figure) or outside the big triangle $v_1v_2v_3v_1$. In all cases, if we wish to connect the vertex v_5 to all the other vertices, it has to be connected with a vertex placed on the opposite side of one of 4 Jordan curves formed by the triangles mentioned above. By Jordan's Theorem the line connecting v_5 to one of the other vertices must intersect one of the triangles and hence it is not possible to draw K_5 in the plane. □

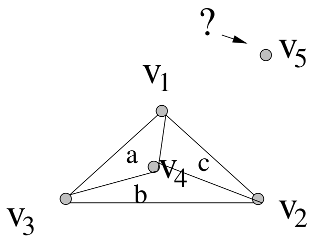

Figure 18.3: The proof of Lemma 136.

Exercise 137 *Show that the graph $K_{3,3}$ (the complete bipartite graph with three vertices on each side) is non-planar.*

One might consider embeddings of graphs on other surfaces. For example, the graph K_5 is embeddable on the torus: we can use the hole in the torus to avoid intersecting edges. See Figure 18.4. The number of holes in the surface is called the *genus* of the surface. The theory we shall developed here for planar graphs has been extended also to higher genus graphs, where surprisingly, some of the results are easier.

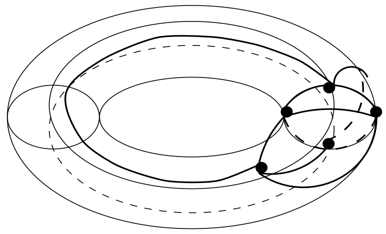

Figure 18.4:

Sometimes it might be more convenient to consider embeddings of graphs on the sphere instead in the plane, because Jordan curves on the sphere have the property that both their inside and its outside regions are bounded, in contrast to Jordan curves in the plane where the outside is unbounded. The following theorem about graphs embedded on the sphere is sometimes useful.

Theorem 138 *A graph G has a planar embedding if and only if it can be embedded on a sphere.*

Proof: Consider the sphere $x^2+y^2+z^2=1$ and the plane $z=-1$. Then, the stereographic projection of the sphere minus the pole $(x,y,z)=(0,0,1)$ on the plane is a bijective function. Thus, if the embedding of the graph on the sphere does not put any vertex or edge on the pole, then the stereographic projection of the vertices and edges in the embedding on the sphere is well defined and continuous curves on the sphere are also continuous in the plane. That tell us that edges do not intersect on the sphere if and only if they do not intersect in the plane. □

From the theorem above we immediately get the following corollary.

Corollary 139 *Given a planar graph G and a vertex v there is a planar embedding of G with v on the unbounded region.*

Proof: To see that, consider the embedding of the graph on the sphere. Rotate the sphere so that v is in the region containing the pole $(0,0,1)$ and project the embedding to the plane. It is clear that v is in the unbounded region, as the stereographic projection of points near the pole are in the unbounded region. □

18.1 The dual graph

The *faces* of a *plane* graph are the connected regions that are left when edges and vertices are removed from the embedding of the graph in the plane. We will denote the set of faces by F. Two faces f_1 and f_2 are *adjacent* if there is an edge e that separates them. Such an edge is said to be *incident* to the two faces. The *degree* of a face $d(f)$ is defined to be the number of edges incident to the face. If an edge touches the same face on both its sides it contributes 2 to the face's degree. See Figure 18.5. It is important to remember that the faces are defined for a plane graph and they depend on the particular embedding of a planar graph.

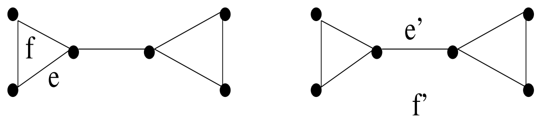

Figure 18.5: Degree of a face: In the left figure the edge e contribute 1 to $d(f)$. In the rigth the edge e' contribute 2 to $d(f')$.

Let $G=(V,E,F)$ be a plane graph. The *dual* graph of G, denoted G^*, is defined as follows: $G^*=(F,E',F')$; for each face in G there is a vertex in G^*; if $e\in E$ is incident to the faces $f_1,f_2\in F$, then $e'=(f_1,f_2)$ is an edge in G^*. The dual G^* of a plane graph is also planar. In fact, there is a natural way to embed it in the plane. For each face of G, draw a vertex of G^* and for each edge e of G draw an edge of G^* that crosses e and connects vertices in G^* the correspond to the faces on the two sides of e. See Figure 18.6.

It is important to remember that the dual of a graph is defined for the *plane* graph. That is, for the specific drawing and not for the planar graph. Just look at the two embeddings (of the same graph) in Figure 18.7. They give rise to two different dual graphs. It is easy to see

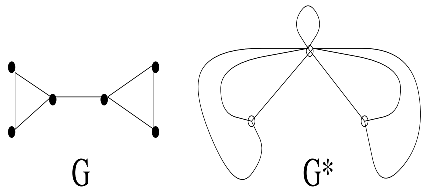

Figure 18.6: A plane graph and its dual.

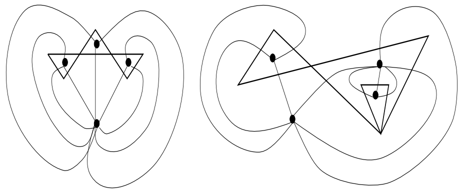

Figure 18.7: Two different embeddings and their duals.

that dual graphs are not isomorphic as the one on the left has three vertices of degree 3 and the one on the right only 2.

It should be easy to convince yourself of the validity of the following claim.

Claim 140 *If G is a connected plane graph then* $(G^*)^* \cong G$

Figure 18.8 shows that if and graph is not connected then the claim above is false.

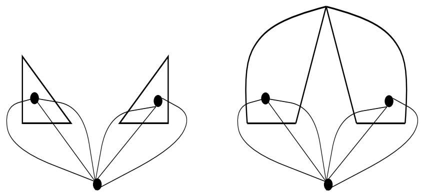

Figure 18.8: An unconnnected graph, its duals and its dual's dual.

Lemma 141 *Let $G = (V, E, F)$ be a plane graph, then*

$$\sum_{f \in F} d(f) = 2 \cdot |E| = 2 \cdot |E^*|.$$

Proof: By construction $|E| = |E^*|$ and by Theorem 55, $\sum_{f \in F} d(f) = 2 \cdot |E^*|$. □

18.2 Euler's Formula

The following simple formula relates the number of faces, vertices and edges of a planar graph.

Theorem 142 (Euler's Formula) *If G is a connected planar graph then*

$$|V| - |E| + |F| = 2.$$

Proof: Let $G = (V, E, F)$ be an arbitrary embedding of the planar graph G. The proof is by induction in the number of faces. If $|F| = 1$, then G must be a tree and we know by Theorem 61, that $|E| = |V| - 1$. Therefore,

$$|V| - |E| + |F| = |V| - (|V| - 1) + 1 = 2.$$

Assume that the formula holds for connected planar graphs with $|F| - 1$ faces. Let G be our plane graph with $|F|$ faces. As $|F| \geq 2$, there must be a cycle in G. Let $e \in E$ be an edge that is not a cut edge. Then, $|E(G-e)| = |E(G)| - 1$, $|V(G-e)| = |V(G)|$ and $|F(G-e)| = |F(G)| - 1$ and we can use induction hypothesis for $G - e$. We get that:

$$|V(G)| - |E(G)| + |F(G)| = |V(G-e)| - (|E(G-e)| + 1) + |F(G-e)| + 1 = 2,$$

establishing the claim. □

Exercise 143 *What has to be changed in Euler's Formula if the planar graph is not connected, but has a fixed number c of connected components?*

Corollary 144 *All embeddings of a connected planar graph G have the same number of faces.*

Proof: Immediate by Euler's Formula. □

Corollary 145 *Let $G = (V, E)$ be a simple planar graph with $|V| \geq 3$ vertices. Then*

$$|E| \leq 3 \cdot |V| - 6.$$

This proves that a planar graph can not have too many edges.

Proof: As a cycle in a simple graph must have at least length 3, we get by Lemma 141

$$3 \cdot |F| \leq \sum_{f \in F} d(f) = 2 \cdot |E|$$
$$\Downarrow$$
$$|F| \leq \frac{2}{3} \cdot |E|.$$

Recall that the claim was about connected graphs, but an upper bound on the number of faces also holds for unconnected graphs (just add edges to connect the graph). From Euler's Formula and the inequality above we now get that

$$\begin{array}{rcl} 2 & = & |V| - |E| + |F| \\ & \leq & |V| - |E| + \frac{2}{3} \cdot |E| \\ & = & |V| - \frac{1}{3} \cdot |E| \\ \Updownarrow & & \\ |E| & \leq & 3 \cdot |V| - 6, \end{array}$$

and the proof is complete. □

Corollary 146 *If $G = (V, E)$ is a simple planar graph then $\delta(G) \leq 5$*

Proof: By Corollary 145,

$$\delta \cdot |V| \leq \sum_{v \in V} d(v) = 2 \cdot |E| \leq 6 \cdot |V| - 12$$
$$\Downarrow$$
$$\delta \leq 6 - \frac{12}{|V|}.$$

And since δ is an integer we get that $\delta \leq 5$. □

We have seen before that K_5 and $K_{3,3}$ are non-planar. We can prove this also by using Euler's Formula.

Corollary 147 *K_5 is non-planar.*

Proof: We know that K_5 is a simple connected graph and that

$$|E(K_5)| = \binom{5}{2} = 10 \quad \text{and} \quad |V(K_5)| = 5.$$

Assuming that it is planar, we can use Corollary 145 and get that

$$10 = |E(K_5)| \leq 3 \cdot 5 - 6 = 9,$$

reaching a contradiction. □

Corollary 148 *$K_{3,3}$ is non planar.*

Proof: Again, we know that $K_{3,3}$ is a simple connected graph and that

$$|E(K_{3,3})| = 9 \quad \text{and} \quad |V(K_{3,3})| = 6.$$

Since $K_{3,3}$ is bipartite, every cycle must have length at least 4 and $d(f) \geq 4$, for every $f \in F$. Now, assuming it is planar, we can use Lemma 141 to get that

$$\begin{aligned} 4 \cdot |F(K_{3,3})| &\leq \sum_{f \in F(K_{3,3})} d(f) = 2 \cdot |E(K_{3,3})| \\ \Downarrow \qquad & \\ |F(K_{3,3})| &\leq \frac{1}{2} \cdot |E(K_{3,3})|. \end{aligned}$$

Using Euler's Formula

$$\begin{aligned} 2 &= |V(K_{3,3})| - |E(K_{3,3})| + |F(K_{3,3})| \\ &\leq |V(K_{3,3})| - |E(K_{3,3})| + \frac{1}{2} \cdot |E(K_{3,3})| \\ &= |V(K_{3,3})| - \frac{1}{2} \cdot |E(K_{3,3})| \\ &= 6 - 4\frac{1}{2} = 1\frac{1}{2}, \end{aligned}$$

reaching a contradiction. □

18.3 Platonic solids

Platonic solids are 3-dimensional polyhedral solids that are highly symmetric: each face and each corner look the same. There exists five platonic solids (their duals are in parenthesis): the tetrahedron (self dual), the cube (octahedron), the icosahedron (dodecahedron), the octahedron (the cube) and the dodecahedron (icosahedron). We prove next that there are no other platonic solids.

Theorem 149 *There exists only 5 platonic solids.*

Proof: For platonic solids there exist the corresponding platonic planar graphs where the corners of the solids are the vertices of the graphs, the faces of the solids are the faces of the graphs and the intersection line of every two faces are the edges of the graphs. The degree of each vertex of the graph has to be the same and at least 3. The degree of each face has also to be the same and also at least 3. Denote by $d \geq 3$ the degree of the vertices and by $p \geq 3$ the degree of the faces. Let $n = \#$ vertices, $m = \#$ edges, $f = \#$ faces. We can immediately establish the following equality.

$$dn = 2m = pf$$

By Euler's Formula we get

$$\begin{array}{lrcl} & n - m + f & = & 2 \\ \Updownarrow & & & \\ & pdn - pdm + pdf & = & 2pd \\ \Updownarrow & & & \\ & 2pm - pdm + 2dm & = & 2pd \\ \Updownarrow & & & \\ & m(2p - pd + 2d) & = & 2pd \\ \Updownarrow & & & \\ & m & = & \frac{2pd}{2p-pd+2d}. \end{array}$$

Since $2pd$ is an integer and m must be a positive integer, the denominator $2p - pd + 2d$ must be a positive integer that divides $2pd$. The only feasible solutions, given in Figure 18.9, are the five platonic solids mentioned above. □

p degree of faces	d degree of vertices	n is # vertices	m is # edges	f is # faces	name of platonic solid
3	3	4	6	4	tetrahedron
3	4	8	12	6	cube
4	3	6	12	8	octahedron
3	5	12	30	20	icosahedron
5	3	20	30	12	dodecahedron

Figure 18.9: The platonic solides.

18.4 Kuratowski's Theorem

Kuratowski's Theorem gives a necessary and sufficient condition for a graph to be planar. Because of time constraints, we are not going to prove it in these lectures and we shall only give the statement of the theorem. Before we can do that, we will need the following facts and definitions.

Fact 150 *Let G be a planar graph and let H be a subgraph of G, then H is planar.*

Proof: Given an embedding of G, start removing the edges and vertices until left with an embedding of H. □

Before we are able to carry on, we need to introduce the concept of a *subdivision*. A *subdivison* of a graph G is a graph H obtained by repeatedly splitting edges in G by placing vertices in their middle. See Figure 18.10. The following fact is immediate.

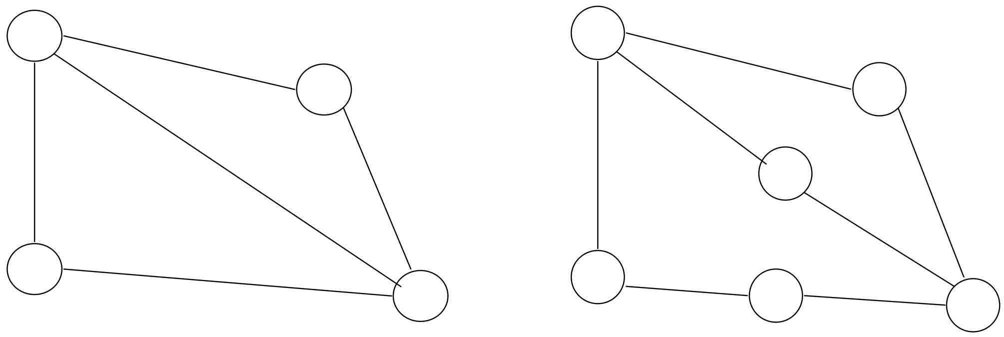

Figure 18.10: A graph and a subdivision.

Fact 151 *Let H be a subdivision of G. Then H is planar if and only if G is planar.*

We are now ready to state the famous Kuratowski's Theorem, which explains our preoccupation with the graphs K_5 and $K_{3,3}$.

Theorem 152 *A graph G is planar if and only if it contains no subgraph isomorphic to a subdivision of K_5 or $K_{3,3}$.*

Notice that one direction in the theorem above easily follows from Facts 150 and 151.

18.5 Colouring planar graphs

The problem of colouring the faces of a planar graphs, so that no two adjacent faces are coloured with the same colour, originates from the need to colour maps in such a way that no two adjacent countries are given the same colour. It was conjectured at the end of the last century that 4 colours always suffice to colour a map, but the proof of this seemingly innocent claim was given only in 1980. (For a long while at the turn of the century, a bogus proof was believed to be correct.) The proof of the 4-colour theorem, however, is too complicated to be presented in these lectures. We shall only prove that 5 colours suffice to colour any map.

If G is a bridgeless planar graph, then G^* is also planar and colouring the faces of G corresponds exactly to colouring the vertices of G^*. Since $(G^*)^* = G$ when G is connected, we can talk only about colouring the vertices of G^*.

Before we prove the 5-colour Theorem, we prove the following easier theorem.

Theorem 153 *Let $G = (V, E)$ be simple planar graphs. Then, $\chi(G) \leq 6$*

Proof: Assume on the contrary that there exist a simple planar graph G such that $\chi(G) > 6$ and let G be a minimal such graph. By Corollary 145, we know that $\delta(G) \leq 5$. Let $v \in V$ be a vertex of G with $d(v) \leq 5$. But then, by the minimality of G, $G - v$ can be coloured with 6 colours and v can be coloured with one of these 6 colours that are not colouring any of the 5 vertices adjacent to v, obtaining a proper 6 colouring of G, in contradiction to the assumption $\chi(G) > 6$. □

Theorem 154 *Let $G = (V, E)$ be a simple planar graph. Then $\chi(G) \leq 5$.*

Proof: Assume on the contrary that there exist a simple planar graph G such that $\chi(G) > 5$ and let G be a minimal such graph example. Pick a vertex $v \in V$, such that $d(v) = 5$. By Corollary 145, there exists a vertex v with $d(v) \leq 5$ and if $d(v) < 5$, then the same proof of Theorem 153 proves that G is 5 colourable.

Take a planar embedding of G and let $v_1, \ldots, v_5$ be the vertices adjacent to v, ordered clockwise. See Figure 18.11. If the vertices $v_1, \ldots, v_5$ are not all coloured with different colours, then again, going back to the proof of Theorem 153, we get that G is 5 colourable.

By the minimality of G, we have that $\chi(G - v) = 5$. Colour $G - v$ with 5 colours and let $V_1, V_2, ..., V_5$ be the subsets of vertices that are coloured with colour $1, 2, \ldots, 5$, respectively. Let $G[V_i \cup V_j]$ be the subgraph of G spanned by the vertices coloured i and j. If v_i and v_j are in different connected components of $G[V_i \cup V_j]$, then we can swap the colours i and j in the component of v_i (or v_j but not both) and obtain a proper colouring in which the neighbours of v are coloured only with 4 different colours. What brings us back again to the proof of Theorem 153.

So, v_i and v_j must be in the same connected component of $G[V_i \cup V_j]$, for all colours i, j. But then, v_1 and v_3 must be connected by a path P_1 with vertices coloured only with colours 1 and 3 and v_2 and v_4 must be connected by a path P_2 with vertices coloured only with colours 2 and 4. See Figure 18.11. Clearly P_1 must intersect P_2. Since G is planar these paths might intersect only at some vertex w. But since w is on the path P_1 it must be coloured either 1 or 3 and since w is also on P_2 it must be coloured either 2 or 4, reaching a contradiction. □

Exercise 155 *Try to prove the 4-colour theorem by repeating the argument in the proof above with G being a minimal graph that is not 5 colourable. What goes wrong?*

Exercise 156 *Give an example of a simple planar graph with $\delta = 5$.*

As mentioned above, it was only in 1980 that the 4-colour theorem was proved. The proof is very long and involved computer verification of several hundred cases.

Theorem 157 (Appel-Haken) *Let $G = (V, E)$ be a simple planar graph. Then $\chi(G) \leq 4$.*

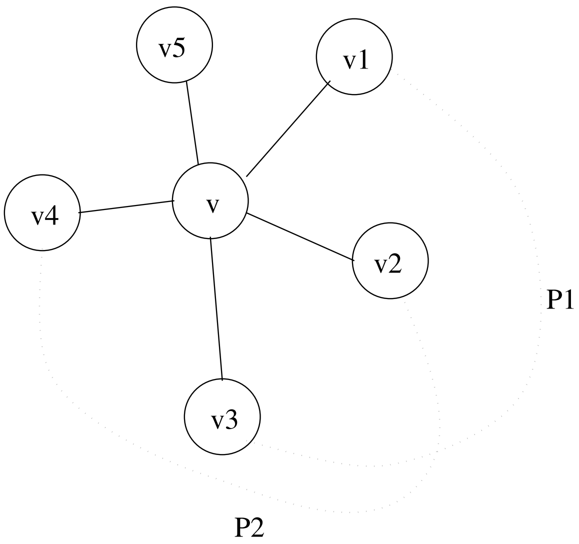

Figure 18.11: The vertex v and its neighbours.

Chapter 19

Problems

The problems below are the same problems given in class. However, their numbers have been permuted.

1. Show that if $G = (V, E)$ is a simple bipartite graph then $|E| \leq n^2/4$.

2. Show that if a tree has a vertex of degree k, then it has at least k leaves (vertices of degree 1).

3. The center of $G = (V, E)$ is a vertex u such that $\max_{v \in V} d(u, v)$ is the smallest possible. Show that a tree has exactly one center or two adjacent centers.

4. A simple graph G is *self complementary* if it is isomorphic to its complement graph G^c. Show that if G is self complementary then $n = 4k$ or $n = 4k + 1$. Construct self complementary graphs for every such n.

5. Let A be the adjacency matrix of some graph G. Show that the number of walks in G from v_i to v_j of length k is the (i, j) entry of A^k.

6. Show that the sequence $d_1, \ldots, d_n$ is the degree sequence of some graph if and only if d_i are non-negative integers and $d_1 + \cdots + d_n$ is even.

7. Prove that the non-negative integer sequence $0 \leq d_1 \leq \cdots \leq d_n$ is the degree sequence of some graph without loops if and only if $d_1 + \cdots + d_n$ is even and $d_n \leq d_1 + \cdots + d_{n-1}$.

8. Prove that the non-negative integer sequence $0 \leq d_1 \leq \cdots \leq d_n$ is the degree sequence of some simple graph if and only if $d_1 + \cdots + d_n$ is even and

$$\sum_{i=n-k+1}^{n} d_i \leq k(k-1) + \sum_{i=1}^{n-k} \min(d_i, k) \quad \text{for} \quad 1 \leq k \leq n.$$

   Prove that $\{d_i\}$ is the degree sequence of a simple connected graph if and only if it satisfies the condition above and $d_1 > 0$ and $d_1 + \cdots + d_n \geq 2(n-1)$.

9. Show that any loopless graph $G = (V, E)$ contains a bipartite spanning subgraph $H = (V, E')$ with $|E'| \geq |E|/2$.

10. Show that if G is simple and $\delta \geq k$, then G has a path of length k.

11. Show that if G is simple and 3-regular then $\kappa(G) = \kappa'(G)$.

12. Show that for any integers $0 < l \leq h \leq g$, there exists a simple graph with $\kappa = l$, $\kappa' = h$ and $\delta = g$.

13. Construct for every $0 \leq k < n$ an n-node k-connected graph with the smallest possible number of edges.

14. Show that a digraph is strongly connected if and only if there is an edge leaving each subset of V.

15. Show that any simple graph with $m > \binom{n-1}{2} + 1$, $n \geq 3$, is Hamiltonian.

16. Prove that for n odd, the edges of K_n can be partitioned into $\frac{n-1}{2}$ Hamiltonian cycles and for n even into $\frac{n}{2}$ Hamiltoniam paths.

17. Let $G = (V, E)$ be a simple graph with $\delta \geq \frac{n+1}{2}$. Then for any edge $e \in E$, there exists a Hamiltonian cycle containing e.

18. An algorithm to find a maximum matching in a graph can be designed by following Berge's Theorem: Starting with an empty matching try to find an augmenting path and increasing the size of the current matching. Consider the following variation on this algorithm. Let l be the length of the shortest M-augmenting path. Find a maximal set of M-augmenting paths of length l and use these paths to increase the size of the current matching.

    Prove that the length l of the shortest augmenting paths increases in each iteration and that the total number of iterations is at most $O(\sqrt{|V|})$.

    This is the basis of the fastest known maximum matching algorithms.

19. Prove that in any bipartite graph $G = (X \cup Y, E)$, the size of a maximum matching is:

$$|X| - \max_{S \subseteq X}(|S| - |N(S)|)$$

20. For each $k > 1$, give an example of a simpe k-regular graph that does not contain a perfect matching.

21. Let G be a $2k$+1-regular $2k$-edge-connected graph. Prove that G has a perfect matching.

22. Let $\alpha(G)$ be the size of the maximum independent set in a graph $G = (V, E)$ and $\Delta(G)$ be the maximum degree of a vertex in G. Prove that $\alpha(G) \geq |V|/(\Delta(G) + 1)$.

23. Prove that if G is a (not necessarily simple graph) with $\Delta = 3$, then $\chi' \leq 4$.

24. Prove that if G is a non-trivial simple regular graph with an odd number of vertices then $\chi' = \Delta + 1$.

25. Let G be a bipartite graph. Prove that it is possible to edge colour G with δ colours such that every colour is present at every vertex.

26. Ramsey's theorem can be stated as follows: if the edges of a simple graph with $n \geq r(k,l)$ vertices are coloured with two colours, then there is a monochromatic K_k or a monochromatic K_l. Consider the following generalization. Let $r(k_1,\ldots,k_h)$ be such that for any complete graph on $n \geq r(k_1,\ldots,k_h)$ vertices whose edges are coloured with h colours there exists some monochromatic complete subgraph K_{k_i}. Prove that:

$$\begin{aligned}&r(k_1,k_2,\ldots,k_h) \leq \\ &\quad r(k_1-1,k_2,\ldots,k_h)+r(k_1,k_2-1,\ldots,k_h)+\cdots+r(k_1,k_2,\ldots,k_h-1).\end{aligned}$$

27. Show that any loopless graph $G=(V,E)$ contains an independent set with $\frac{|V|^2}{|V|+2|E|}$ vertices.

28. Using the independent set bound above, show that a parallel comparison algorithm that makes n comparisons in each round must take at least $\Omega(\log\log n)$ rounds to find the maximum of n elements. (Hint: use an adversary argument. Consider initially all n elements as candidates to be the maximum. Draw a graph with the candidates being the vertex set and the comparisons being the edge set. When a round of comparisons is made, an independent set of candidates can be chosen as the candidates for the next round.) Give a comparison model algorithm with matching bound. What are the bounds if the number of comparisons in each round is p and the number of elements is n.

29. Prove that the faces of a plane graph G are 2-colourable if and only if G is Eulerian.

30. Let $G=(V,E)$ be a planar graph. Prove that it has an induced bipartite subgraph $H=(V',E')$ such that $|V'| \geq |V|/2$.

Solutions

Most of the solutions below have been provided by the students participating the course. We have sometimes edited the solutions and simplified them.

Solution 1: The proof is a consequence of Lemma 121.

Solution 2 (Theis Rauhe): Let $d_1,\ldots,d_n$ be the degrees of the vertices in a tree T. Since a tree is connected $d_i \geq 1$. Suppose that T has a vertex of degree k, but fewer than k leaves. Let h be the number of leaves in T. We can assume that the leaves are the vertices with degrees $d_1=d_2=\cdots=d_h$, that $d_{h+1}=k$ and that $d_i \geq 2$ for $i \geq h+2$. But then,

$$\begin{aligned}\sum_i d_i &= d_1+\ldots+d_h+d_{h+1}+d_{h+2}+\ldots+d_n \\ &= h+k+d_{h+2}+\cdots+d_n \\ &\geq h+k+2(n-(h+1)) \\ &\geq 2n-1.\end{aligned}$$

But this contradicts the fact that $\sum d_i = 2n-2$ by Theorems 61 and 55.

Solution 3 (Thomas Hildebrandt): Since V is finite there must always be at least one center. If $|V| = 1$, then the one vertex is the center, trivially. If $|V| \geq 2$, assume that v_1 is a center. Let u be a vertex with $m = d(v_1, u)$ largest possible and let $P : v_1 \to w \rightsquigarrow u$ be the path from v_1 to u. There must be at least one vertex z different from u, such that $d(v_1, z) \geq m - 1$ and w is *not* on the path P' from v_1 to z, for otherwise $max_{v \in V} d(w, v) < max_{v \in V} d(u, v)$ contradicting that v_1 is a center. I.e. we have a situation like:

$$\bullet_z \rightsquigarrow^{\geq m-1} \cdots \rightsquigarrow \bullet_{v1} \to \bullet_w \rightsquigarrow^{m-1} \cdots \rightsquigarrow \bullet_u$$

Now assume we have another center v_2. Let $x \in P \cup P'$ be a vertex closest to v_2. Since $x \in P \cup P'$ either $d(x, z) \geq m$ or $d(x, u) \geq m$. Without loss of generality assume $d(x, z) \geq m$. Since G is acyclic we then get $d(v_2, z) = d(v_2, x) + d(x, z) \geq d(v_2, x) + m$, which implies that $d(v_2, x) = 0$, or that $v_2 \in P \cup P'$. So all possible centers are on the paths P and P'.

Except for w, no other vertex on $P \cup P'$ can be a center (it's either too far from z or too far from u). Thus, there cannot be more than these two centers, and if there are two, they are adjacent.

Solution 4: Let G be a self-complementary simple graph. Then, the number of edges m in G and in G^c has to be equal to $m = \binom{n}{2}/2 = n(n-1)/4$. But since this is an integer, either n or $n - 1$ must be divisible by 4 and $n = 4k$ or $n = 4k + 1$, for some integer k.

We give next examples of self-complementary graphs with 4 and 5 vertices.

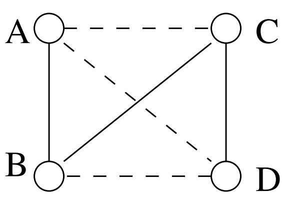

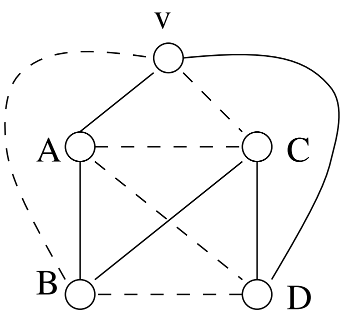

Self-complementary graphs with $4k$ and $4k+1$ vertices can be obtained by replacing A and D by the complete graph K_k and B and C by the complement of K_k (which has k nodes and no edges), while v remains a single node. The edges are replaced by collections of edges that connect all pairs of vertices in the corresponding sets. (I.e. all vertices in A with all vertices in B, and so on.) Other replacements for A, B, C, D and for the edges are also possible.

Solution 5 (Thomas Hildebrandt): By induction on k. For $k = 0$ we get $A^0 = I$ which agrees with the fact that there is exactly one walk from v_i to v_j of length 0 if and only if $i = j$. For $k = 1$ we get $A^1 = A$. By definition the (i, j)'th entry of the adjacency matrix is the number of walks from v_i to v_j of length 1.

Assume the statement hold for walks of length at most k. Any walk of length $k + 1$ from v_i to v_j can be partitioned into the first edge to a vertex v_l and the rest of the walk of length k from v_l to v_j, i.e.

$$\bullet_i \to^1 \quad \bullet_l \rightsquigarrow^k \cdots \rightsquigarrow \quad \bullet_j$$

So for each possible first edge to v_l, the number of walks of length $k + 1$ is the number of walks of length 1 between v_i and v_l multiplied by the number of walks of length k between v_l

and v_j, which by induction is $A_{il}A^k_{lj}$. If we add this together for all possible first edges, we get: $\sum_{l=1}^n A_{il}A^k_{lj}$, which exactly is the (i,j)'th entry of A^{k+1}.

Solution 6 (Theis Rauhe): By Theorem 55 and the definitions, the condition is necessary. Suppose that $d_1, \ldots, d_n$ are non-negative integers and $\sum_i d_i$ is even. The proof is constructive by induction on n.

$n = 1$: Since d_1 must be even we construct a graph consisting of one vertex with $\frac{d_1}{2}$ loops by which it has degree d_1.

$n = 2$: There are two cases: if both d_1 and d_2 are even, construct a graph with two vertices v_1, v_2 with $\frac{d_1}{2}$ and $\frac{d_2}{2}$ loops, respectively. If both d_1 and d_2 are odd let v_1 and v_2 have $\frac{d_1-1}{2}$ and $\frac{d_2-1}{2}$ loops, respectively, and by connecting v_1 and v_2 by a single edge, v_1 and v_2 get degree d_1 and d_2.

$n > 2$: Suppose the claim holds for $1 \leq k < n$. Let $\{d_i, d_j\}$ be a pair such that $d_i + d_j$ is even (always possible since $n > 2$). Construct by induction hypotheses a separate graph for d_i, d_j and for $d_1, \ldots, d_{i-1}, d_{i+1}, \ldots, d_{j-1}, d_{j+1}, \ldots d_n$.

Solution 7 (Lone Overgaard): Assume that $0 \leq d_1 \leq \cdots \leq d_n$ is the degree sequence of some loopless graph. By Theorem 55 the sum of the degrees is even. Since each edge that has one end incident to the vertex v_n must have its other end incident to another vertex, we get that $d_n \leq d_1 + \cdots + d_{n-1}$.

Now assume that $d_1 + d_2 + \cdots + d_n$ is even and $d_n \leq d_1 + \cdots + d_{n-1}$. By the previous exercise $d_1, \ldots, d_n$ is the degree sequence of some graph (with loops). Let G be such a graph with the smallest number of loops. G can not have two vertices that are incident to loops for otherwise the two loops can be eliminated by puting one of their ends on the other vertex, without changing the degrees and contradicting the minimality of G. Assume that there is a loop at a vertex v. If G has an edge uw, $u, w \neq v$, then the loop can eliminated by connecting one of its ends to u and replacing the edge uw by vw, without changing the degrees and again contradicting the minimality of G. Therefore, all edges in G must be connecting v to the other vertices or are loops with both ends incident to v, contradicting that $d_n \leq d_1 + \cdots + d_{n-1}$.

Solution 8 (Peer Duschner): Simple graphs. Suppose that $\{d_i\}$ is the degree sequence of some simple graph. Consider any subset of k vertices—how big can the degree sum of these vertices be? There can be at most $\binom{k}{2}$ edges between the k vertices contributing $k(k-1)$ to the sum, and the remaining $n - k$ vertices can contribute at most $\min(d_i, k)$ edges, each edge adding 1 to the sum. The inequality follows as a special case.

Let $\{d_i\}$ be a sequence with even sum for which the inequalities hold. We prove the claim by induction on the degree sum. The base clearly holds for sequences with degree sum zero. The induction step will assume that all degrees are non-zero as vertices of degree zero can be added at any time.

If the degrees are all the same, then the graph is regular and we have seen before how such a graph can be constructed. Let l be the smallest index such that $d_l = d_n$. We will show that the degree sequence

$$d_1 - 1 \leq \cdots \leq d_{l-1} \leq d_l - 1 < d_{l+1} = \cdots = d_n,$$

which clearly has also an even sum, satisfies the inequalities. Then, by the induction hypothesis, there is a simple graph with this new degree sequence (of smaller sum), and we can obtain a simple graph with our original degree sequence $\{d_i\}$ by connecting the vertices of degrees

$d_1 - 1$ and $d_l - 1$, if they are not connected. If these vertices are connected, call them u and v respectively, and let w be a vertex of non-zero degree that is not connected to v. Since the degree of w is larger than the degree of u, there must be a vertex x that is connected to w but not to u. Remove the edges wx and uv and connect ux and vw.

Let us now prove the inequalities for the new degree sequence, denoted $\{d_i'\}$. Assume that some inequality does not hold and let k be the smallest index for which

$$\sum_{i=n-k+1}^{n} d_i' > k(k-1) + \sum_{i=1}^{n-k} \min(d_i', k).$$

We must have that $l + 1 \leq n - k$, for otherwise the inequality in the original degrees also does not hold, since we took away at least 1 from the left side and at most 1 from the right side. If $d_n \leq \sum_{i=1}^{n} \min(d_i, 1)$, then trivially $d_n' \leq \sum_{i=1}^{n} \min(d_i', 1)$, implying that $k \geq 2$. If the equality

$$\sum_{i=n-k'+1}^{n} d_i' \leq k'(k'-1) + \sum_{i=1}^{n-k'} \min(d_i', k')$$

holds for $k' < k$, such that $d_{n-k'}' \leq k'$, it must also hold for $k'+1$. We therefore must conclude that $d_n \geq k$. But if $d_n = k$, then

$$k^2 = \sum_{i=n-k+1}^{n} d_i' \leq k(k-1) + \sum_{i=1}^{n-k} \min(d_i', k),$$

holds clearly if $n - k > 2$ and also if $n - k = 2$, for otherwise the degree sum is odd. If $d_n > k$, then, we show that

$$\sum_{i=n-k+1}^{n} d_i < k(k-1) + \sum_{i=1}^{n-k} \min(d_i, k),$$

and therefor,

$$\sum_{i=n-k+1}^{n} d_i' \leq k(k-1) + \sum_{i=1}^{n-k} \min(d_i', k).$$

Assume that this is not the case and let h be the smallest index such that $d_h \geq k + 1$. Then,

$$\begin{aligned}
\sum_{i=n-k+1}^{n} d_i &= k(k-1) + \sum_{i=1}^{n-k} \min(d_i, k) \\
&= k(k-1) + \sum_{i=1}^{h-1} d_i + k(n-k-h+1) \\
&= k(n-h) + \sum_{i=1}^{h-1} d_i,
\end{aligned}$$

and

$$\begin{aligned}
\sum_{i=n-k}^{n} d_i &= k(n-h) + \sum_{i=1}^{h-1} d_i + d_{n-k} \\
&> (k+1)k + \sum_{i=1}^{n-k-1} \min(d_i, k+1)
\end{aligned}$$

$$
\begin{aligned}
&= (k+1)k + \sum_{i=1}^{h-1} d_i + (k+1)(n-k-h) \\
&= (n-h+1)k + \sum_{i=1}^{h-1} d_i,
\end{aligned}
$$

since $d_n = d_{n-k} > k$, contradicting the originial inequality in the sequence $\{d_i\}$ for $k+1$.
Simple connected graphs. By Theorem 55 and Corollary 67, in a connected graph the degree sum is at least $2(n-1)$ and clearly the degree of each vertex is at least one.

Let $\{d_i\}_{1\leq i\leq n}$ be a degree sequence satisfying the conditions. As we proved above, there exists a simple graph with this degree sequence. Let G be such a graph with the smallest number of connected components. Since the degree of each vertex is at least one, we know that each component must have at least two vertices. If G is not connected, since the degree sum is at least $2(n-1)$, we know that there is at least one component that contains a cycle. Let uv be an edge in G on this cycle and let xy be an edge in a different component of G. By removing the edges uv and xy and replacing them by ux and vy, we have created a single connected component containing all the vertices in the components of uv and xy, and thus, we have reduced the number of components without changing the degrees, in contradiction to the minimality of G.

Solution 9 (Theis Rauhe): Let H be a bipartite spanning subgraph of G with maximal number of edges. We claim that for each vertex v, $d_H(v) \geq \frac{1}{2}d_G(v)$. If this isn't the case, v has to be connected with more than $\frac{1}{2}d_G(v)$ vertices in the part containing it. But then we could move it to the other part and increase the number of edges in H contradicting the assumption of H being maximal.

Solution 10 (Søren Skyum): Lets prove a stronger statement by induction: For every vertex v there is a path of length k starting at v. The basis $k=1$ holds since every vertex is connected to at least one other vertex. Assume that the statement holds for any graph with $\delta \geq k-1$. Choose any vertex v. By assumption $d(v) \geq k$. Choose one of v's neighbours u, and remove v from the graph. The new graph G' has $\delta' \geq k-1$. The induction assumption tells us that there exists a path of length $k-1$ in G' starting at u. Therefore, the original graph has a path of length k starting at v.

Solution 11 (Thomas Hildebrandt): Assume G is simple and 3-regular. By Theorem 93, $\kappa(G) \leq \kappa'(G) \leq \delta(G) = 3$. This leaves us with three cases to prove:
Case 1: $\kappa(G) = 3$. Trivially $\kappa'(G) = 3$.
Case 2: $\kappa(G) = 2$. The graph must be separated in one of the three ways sketched in Figure 19.1. It is easy to see that the graph can be disconnected by removing the bold edges, i.e. $\kappa'(G) = 2$.
Case 3: $\kappa(G) = 1$. The graph must be separated as sketched in Figure 19.2. It is easy to see, that the graph can be disconnected by removing the bold edge, i.e. $\kappa'(G) = 1$.

Solution 12: Take two copies A and B of the complete graph K_{g+1}. Choose l vertices $a_1, \ldots, a_l$ in A and h vertices $b_1, \ldots, b_h$ in B and connect b_i to a_i, for $i = 1, \ldots, l$ and each b_j to an arbitrary a_q, for $j = l+1, \ldots, h$. In the resulting graph, clearly $\delta = g$, $\kappa = l$ (the smallest vertex separator is $\{a_1, \ldots, a_l\}$) and $\kappa' = h$ (the smallest cut is (A, B).)

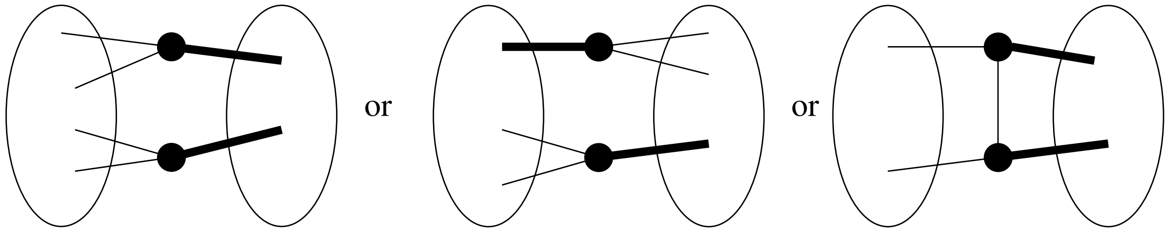

Figure 19.1: Case 2

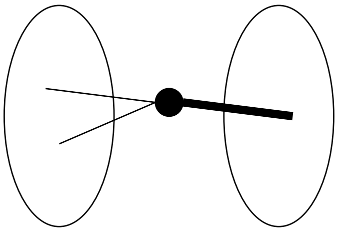

Figure 19.2: Case 3

Solution 13: For $k = 0$, the required graph clearly has no edges. For $k = 1$, by Corollary 67 the required graph must contain at least $n - 1$ edges and any spanning tree would do. Recall that by Example 94, a k-connected graph with n vertices must have at least $kn/2$ edges.

If k or n are even, then $kn/2$ is an integer and the k-regular n-vertex graph constructed in Exercises 57 has exactly $kn/2$ edges. We claim that this graph is indeed k-connected.

We prove this here only in the case where k is even. Assume that the graph has a vertex cut S, $|S| < k$, that separates the vertices i from j. Consider the two sets of vertices on the circle between i and j and the vertices between j and i. One of these sets has fewer than $k/2$ vertices from S, and thus, this set must contain a sequence of vertices that are at most $k/2$ apart and containing i and j. By the construction these vertices are connected, making such a separator impossible.

If both k and n are odd, then kn is odd, and the number of edges must be at least $\lceil kn/2 \rceil$. It is possible to construct a k-connected n-vertex graph with this number of edges by taking the $k-1$-regular n-vertex graph construct in Exercise 57 connecting each vertex number i on the circle to the almost opposite vertex whose number is $i+(n+1)/2$, for $i = 1, \ldots, (n-1)/2$ and connecting vertex number 1 to vertex number $1+(n-1)/2$. Here we also omit the proof that this graph satisfies the requirements.

Solution 14 (Theis Rauhe): Let $G = (V, E)$ be a strongly connected digraph. Suppose that there is a non-trivial set $V' \subset V$ such that no edge leaves V'. Let $v_1 \in V'$ and $u \in V \setminus V'$. Since G is strongly connected there is a directed path from v_1 to u, and some edge on this path must leave V'.

Let $G = (V, E)$ be a digraph such that each non-trivial subset $V' \subset V$ has an edge leaving V' and suppose that G is not strongly connected. That is, there are two vertices v and u, such that there is no path from v to u. Consider the set $V_v = \{w \mid \text{there is a path from } v \text{ to } w\}$. V_v is a non-trivial subset of V since $v \in V_v$ and $u \notin V_v$. By assumption, there is an edge leaving V_v, say (v', w), $v' \in V_v$ and $w \notin V_v$. By the definition of V_v there is a path from v to v' and therefore there is a path from v to w, contradicting that $w \notin V_v$.

Solution 15 (Lone Overgaard): Notice that $m > \binom{n-1}{2} + 1$ means a graph where

$$\sum_{i=1}^{n} d_i = 2m > (n-1)(n-2) + 2 = n^2 - 3n + 4.$$

We shall try to reach a contradiction in Theorem 89 by assuming that there exist a $k < n/2$, such that

$$d_k \le k \quad \text{and} \quad d_{n-k} < n-k.$$

But then,

$$\begin{aligned}\sum_{i=1}^{n} d_i &= \sum_{i=1}^{k} d_i + \sum_{i=k+1}^{n-k} d_i + \sum_{i=n-k+1}^{n} d_i \\ &\le k^2 + (n-k-1)(n-2k) + (n-1)k\end{aligned}$$

Define $f(k) := k^2 + (n-k-1)(n-2k) + (n-1)k$. Since this is a quadratic function with positive leading coefficient,

$$f(k) \le \max(f(1), f(\frac{n}{2}-1)) \le \max(n^2 - 3n + 4, \frac{3}{4}n^2 - \frac{3}{2}n + 2) \quad \text{for} \quad k \in [1, \frac{n}{2} - 1].$$

But $n^2 - 3n + 4 \ge \frac{3}{4}n^2 - \frac{3}{2}n + 2$, if $n \ne 3$, and

$$\sum_{i=1}^{n} d_i \le n^2 - 3n + 4,$$

contradicting the assumption that there was such a k. We conclude that by Theorem 89 the graph is Hamiltonian. If $n = 3$, then the claim trivially holds.

Solution 16: If n is odd and the edges of K_n are partitioned into $\frac{n-1}{2}$ edge disjoint Hamiltonian cycles then by deleting one vertex we obtain K_{n-1} whose edges are partitioned into $\frac{n-1}{2}$ disjoint Hamiltonian paths. Similarly, if n is even and the edges of K_n are partitioned into $\frac{n}{2}$ edge disjoint Hamiltonian paths, then by adding over vertex and connecting it to all n vertices we get K_{n+1} with its edges partitioned into $\frac{n}{2}$ Hamiltonian cycles. So it suffices to prove only one of these partitions.

Assume that n is even and consider the Hamiltonian path $P(0) = 1, n, 2, n-1, 3, n-2, \ldots, n/2, n/2+1$. Define for $i = 0, \ldots, \frac{n}{2} - 1$, the Hamiltonian path $P(i)$ to be the same as $P(0)$ with the vertex numbers shifted by i modulo n. We claim that the paths $P(i)$ are edge disjoint. This should be proved formally, but it can also be easily seen in a drawing:

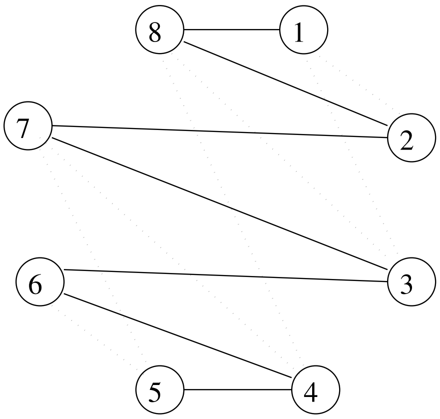

In the special case where n is an odd prime, edge disjoint Hamiltonian cycles can also be constructed by connecting vertex number i to $i+h$, for $h = 1, \ldots, \frac{n-1}{2}$. This however fails to if n is not a prime since we do not get necessarily cycles.

Solution 17 (Søren Skyum): Remove the edge $e = (u, v)$ from the graph. Insert a new vertex w, and connect it to u and v. If this new graph G' has a Hamiltonian cycle, the original graph G must have a Hamiltonian cycle containing e. We order the degrees of the vertices in G': $d_1 \leq d_2 \leq \cdots \leq d_{n'}$, where $n' = n + 1$. We know that $d_1 = 2$ and $d_i \geq \frac{n'}{2}$, for $2 \leq i \leq n'$. Appealing to Theorem 89, we ask if there is a $k < n'/2$, such that $d_k \leq k$ and $d_{n'-k} < n' - k$? $d_1 = 2$, so it is not possible that $k = 1$ and by assumption $d_k \geq n'/2$, establishing that there is no such $k < n'/2$. Since no such k exists, by Theorem 89 G' is Hamiltonian and so is G.

Solution 18 (Thomas Hildebrandt): The proof presented here is essentially from the paper [Hopcroft and Karp, SIAM J. Comput. 1973, An $n^{5/2}$ Algorithm for Maximum Matchings in Bipartite Graphs]

Define $M \oplus P = (M \cup P) \backslash (M \cap P)$, the operation of augmenting a matching M with an augmenting path(s) P. Let S be a maximal set of vertex disjoint M-augmenting paths of minimal length l, and let $M' = M \oplus S$ and s be the number of paths in S. We will now show that the length of any M'-augmenting path is strictly bigger than l.

Let P be an augmenting path in M'. Let $N = M' \oplus P = (M \oplus S) \oplus P$. Then N is a matching of size $|N| = |M| + s + 1$. By the same argument as in the proof of Berge's theorem, namely that $M \cup N$ consists of disjoint MN-alternating paths and even cycles, $M \cup N$ contains $s+1$ disjoint M-augmenting paths, and so does $M \oplus N$. Call these paths $P_1, \ldots, P_{s+1}$. Since $M \oplus N = S \oplus P$, we get $|S \oplus P| \geq \sum_{1 \leq i \leq s+1} |P_i|$ and since $|P_i| \geq l$, we have $|S \oplus P| \geq (s+1) \cdot l$. But we also have the identity $|S \oplus P| = |S| + |P| - 2|S \cap P|$, i.e. $|S \oplus P| = s \cdot l + |P| - 2|S \cap P|$. Hence $|P| \geq l + 2|S \cap P|$. Since S was maximal, if $|P| = l$, P must intersect some of the paths in S, but this gives a contradiction, so we have $|P| > l$.

From the first part above we get that after $\sqrt{|V|}$ steps, the length of the shortest augmenting path $l \geq 2\sqrt{|V|} + 1$. Let M be the matching obtained after $\sqrt{|V|}$ steps and let M^* be a maximum matching. If $m = |M^*| - |M|$, then $M^* \cup M$ must contain at least m vertex disjoint M-augmenting paths. But since these paths are of length at least $\sqrt{|V|}$ we can only have at most $\sqrt{|V|}$ of them, so $m \leq \sqrt{|V|}$. Since each step increases the size of the matching, we can only have $m = \sqrt{|V|}$ more steps. The number of iterations is therefor bounded by $2 \cdot \sqrt{|V|}$.

Solution 19: Let M be a maximum matching and let $X' \subseteq X$ be the set of M-saturated vertices in X. Let

$$h = \max_{S \subseteq X} (|S| - |N(S)|).$$

If we add h vertices on the side of Y and connect them all to all vertices in X, we get a graph in which the conditions of Hall's Theorem are satisfied, so there is a matching saturating all of X. At most h of the vertices in X are matched to these additional vertices, so the rest of the $|X| - h$ vertices are matched within G and $|M| \geq |X| - h$.

Consider the subgraph G' induced by $X' \cup Y$. By Hall's Theorem, for any subset $S \subseteq X'$, $|N_{G'}(S)| \geq |S|$. Therefore, for any $S \subseteq X$, $|N_G(S)| \geq |N_{G'}(S \cap X')| \geq |S \cap X'|$. Thus, if we let S be the subset for which h gets its maximum value,

$$|X| - h = |X| - (|S| - |N_G(S)|) \geq |X| - (|S \cap X'| - |N_{G'}(S \cap X')| + |S \setminus X'|) \geq$$
$$|X| - |S \setminus X'| \geq |X'| = |M|.$$

Solution 20 (Søren Skyum): If k is even the complete graph K_{k+1} doesn't contain a perfect matching, since it has an odd number of vertices. If k is odd an example can be constructed in the following way: Take k copies of K_{k+1}. In each of these remove one edge (u_i, v_i) and add a new vertex w_i. Connect w_i to u_i and v_i. This results in k graphs, where each vertex except one (w_i) has degree k. Add $k-2$ new vertices to the graph. Connect each of these to all the w_i's. The resulting graph is now k-regular. If there is a perfect matching in this graph, the vertices w_i have to be matched outside their original K_{k+1}, since there is an even number of vertices in these. But the k w_i's are only connected to $k-2$ vertices outside, making a perfect match impossible.

Solution 21 (Lone Overgaard): This is a generalizations of Petersen's Theorem. Tutte's Theorem gives us that G has a perfect matching if $odd(G-S) \leq |S|$, for all $S \subseteq V$. Let S be an arbitrary subset of V and let $G_1, \cdots, G_n$ be the odd components of $G - S$. Since G is $2k+1$-regular we know that

$$\sum_{v \in V(G_i)} d(v) = (2k+1)|V(G_i)|$$

and

$$\sum_{v \in S} d(v) = (2k+1)|S|$$

Now let $m_i = \#$ edges with endpoint in G_i and $V - V(G_i)$. As no two odd components G_i and G_j are connected in $G - S$ for $i \neq j$, m_i is the number of edges between the odd component G_i and S. G is $2k$-edge-connected so $m_i \geq 2k$ and even more precisely we can get

$$\begin{aligned} m_i &= \sum_{v \in V(G_i)} d(v) - 2|E(G_i)| \\ &= (2k+1)|V(G_i)| - 2|E(G_i)|, \end{aligned}$$

that is m_i must be odd and therefor $m_i \geq 2k+1$. Combining these results we get

$$\begin{aligned} odd(G-S) &= n \\ &\leq \frac{1}{2k+1} \sum_{i=1}^{n} m_i \\ &= \frac{1}{2k+1} \sum_{v \in S} d(v) \\ &= |S|, \end{aligned}$$

and the number of odd components in $G - S$ for an arbitrary subset S of V is less than or equal to the size of S. By Tutte's Theorem G has a perfect matching.

Solution 22: An independent set can be clearly constructed by repeatedly picking an arbitrary vertex, adding it to the independent set and removing it and all its neighbours from the graph. This step can be repeated at least $|V|/(\Delta+1)$ times resulting in an independent set with at least this number of vertices.

Solution 23 (Søren Skyum): Remove all multiple-edges from G creating a simple graph G'. This new graph can be coloured by 4 colours by Vizing's theorem. Let u and v be two vertices connected by two edges in G. The colouring of G' looks like figure 19.3 (modulo numbering). The extra edge can of course be coloured by the last colour c_4. If two vertices

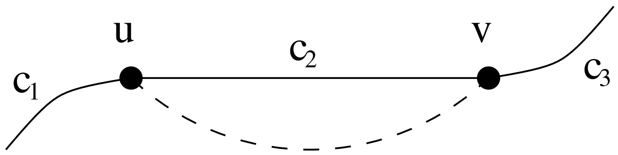

Figure 19.3: The colouring of G'

are connected by three edges, they are not connected to the rest of the graph, and these edges can be coloured by any three colours.

Solution 24: By Vizing's Theorem, $\Delta \leq \chi'(G) \leq \Delta + 1$. The edges coloured with any given colour form a matching and if the number of vertices is odd then in any matching there is at least one vertex that is not matched. Since G is regular, each vertex has degree Δ, but the maximum degree sum in Δ matchings is $\Delta|V| - \Delta$, establishing that at least $\Delta + 1$ colours are necessary to properly colour the edges of G.

Solution 25: The proof is very similar to that of Theorem 110. Let C be an optimal δ-edge-colouring of the bipartite graph G. We claim that $c(v) = \delta$ for all vertices v. Assume on the contrary that there exists a vertex v such that $c(v) < \delta$. Then, at least two edges incident to v have the same colour i and there is some colour j that is not present at v. By Lemma 109, G must have an odd cycle, contradicting that it is a bipartite graph.

Solution 26: The proof is almost identical to that of Theorem 117. Clearly $r(k_1, \ldots, k_h) = 1$ if at least one of the k_i's is 1. The proof is by induction on the k_i's. Let G be the complete graph with

$$r(k_1 - 1, k_2, \ldots, k_h) + r(k_1, k_2 - 1, \ldots, k_h) + \cdots + r(k_1, k_2, \ldots, k_h - 1)$$

vertices and colour the edges of G with h colours. Consider an arbitrary vertex v and let $N_i(v)$ denote the set of vertices that are adjacent to v by edges of colour i. Then, there must exist some colour i, such that,

$$N_i(v) \geq r(k_1, \ldots, k_i - 1, \ldots, k_h).$$

By the induction hypothesis, the subgraph induced by the vertices $N_i(v)$ either contains a monochromatic clique K_{k_j}, for some $j \neq i$, or a monochromatic clique $K_{k_i - 1}$, which together with v forms a monochromatic clique of size k_i.

Solution 27 (Peer Duschner): Without loss of generality assume that $G = (V, E)$ is a simple graph. Let $n := |V|$ and $m := |E|$. Assume that G has an independent set of size h, but not of size $h+1$. Then, G^c has a clique of size h, but no clique of size $h+1$. By Turàn's Theorem,

$$|E(G^c)| = \binom{n}{2} - m \leq \sum_{0\leq i<j<h} \lfloor \frac{n+i}{h} \rfloor \lfloor \frac{n+j}{h} \rfloor.$$

We first try to get rid of the rounding. Let $n = kh + h'$, with $0 \leq h' < h$. Then,

$$\begin{aligned}
\sum_{0\leq i<j<h} \lfloor \frac{n+i}{h} \rfloor \lfloor \frac{n+j}{h} \rfloor &= \sum_{0\leq i<j<h-h'} k^2 + \sum_{0\leq i<h-h'\leq j<h} (k^2+k) + \sum_{h-h'\leq i<j<h} (k^2+2k+1) \\
&= \binom{h}{2} k^2 + h'(h-1)k + \binom{h'}{2} \\
&\leq \binom{h}{2} (k + \frac{h'}{h})^2 = \binom{h}{2} (\frac{n}{h})^2.
\end{aligned}$$

But solving,

$$\frac{n(n-1)}{2} - m = \binom{n}{2} - m \leq \binom{h}{2} (\frac{n}{h})^2 = \frac{h(h-1)}{2} \frac{n^2}{h^2},$$

we get that

$$\frac{n^2}{n+2m} \leq h,$$

which is what we had to prove.

Solution 28 (Thomas Hildebrandt): To find the maximum of n elements we have to compare elements until all but one element have "lost" at least one comparison. We represent the progress of the algorithm by a graph with n vertices that correspond to the input elements and comparisons are represented by edges between the vertices corresponding to the elements being compared. Initially the graph has no edges and all vertices (elements) are candidates to be the maximum. The lower bound proof is given by an adversary argument. The adversary maintain that after step number h, there is a large set I_h of candidates that can still be the maximum. The initial candidates in I_0 are all the n input elements.

Suppose that at step h an algorithm makes n comparisons. If the algorithm was "smart", it would make comparisons only between pairs of elements that are still candidates to be the maximum at the end of the previous step (other comparisons are waisted from the point of view of the adversary.) The adversary maintains that I_h will be an independent set in the graph modeling the computation, and thus, we can assume that all elements in this independent set are larger than all other elements and thus are still candidates for the maximum. But from the previous exercise, we know that there exists always an independent set $I_h \subseteq I_{h-1}$, with at least $\frac{|I_h|^2}{|I_h|+2n}$ vertices. The algorithm can not terminate with the correct answer before $|I_h| = 1$. Let $k_0 = 1$ and $k_{i+1} = \frac{k_i^2}{k_i+2}$, then, $|I_0| = k_0 n$ and

$$|I_{i+1}| \geq \frac{(k_i n)^2}{(k_i n) + 2n} = \frac{k_i^2}{(k_i+2)} n = k_{i+1} n.$$

The adversary can make any algorithm go on until $k_h \leq n$. Instead of calculating with k_i we use the estimate $\tilde{k}$, where $\tilde{k}_1 := \frac{1}{3}$, and for $i \geq 1$, $\tilde{k}_{i+1} := (\tilde{k}_i)^3$, i.e. $\tilde{k}_{i+1} := (\frac{1}{3})^{3^i}$. Now:

$$\begin{array}{rcl} \tilde{k}_{h+1} & > & \frac{1}{n} \\ \Updownarrow & & \\ (\frac{1}{3})^{3^h} & > & \frac{1}{n} \\ \Updownarrow & & \\ \log_3 \log_3 n & > & h \end{array}$$

It is easy to check that $k_i \geq \tilde{k}_i$ for all $i \geq 1$, so from above it follows that it takes at least $\log_3 \log_3 n$ rounds before k_h is below $\frac{1}{n}$, i.e. we have the promised lower bound on $\Omega(\log \log n)$.

In the general case with $p \geq n$ comparisons in each rounds, we get by the same reasoning $k_{i+1} = \frac{k_i^2}{(k_i + \frac{2p}{n})}$. Again we can use an estimate $\tilde{k}$, where $k_i \geq \tilde{k}_i$ for all $i \geq 1$, by defining $\tilde{k}_1 := \frac{1}{1+\frac{2p}{n}}$, and for $i \geq 1$, $\tilde{k}_{i+1} := (\tilde{k}_i)^3$, i.e. $\tilde{k}_{i+1} := (\frac{1}{1+\frac{2p}{n}})^{3^i}$. From this we get the lower bound:

$$\log(\frac{\log n}{\log(1 + \frac{2p}{n})})$$

We also have the lower bound $\frac{n}{p}$, if $p < n$, since this number of rounds are necessary to compare each element at least once. This combines to the lower bounds

$$\Omega(\log \log_{1+2p/n} n + \frac{n}{p}).$$

We now devise an algorithm for finding the maximum. In the case with n processors, the algorithm partitions the elements in $\sqrt{n}$ groups and computes maximum of these in parallel by $\sqrt{n}$ recursive calls, each using $\sqrt{n}$ processors. When these calls return we have $\sqrt{n}$ candidates to be the maximum, which can be done in constant time with n processors by comparing all pairs of elements. So an upper bound on the running time is the depth of the recursion, which is the number of times one must take the square root of n to get below 2:

$$\begin{array}{rcl} n^{2^{-i}} & \leq & 2 \\ \Updownarrow & & \\ 2^{-i} \log n & \leq & 1 \\ \Updownarrow & & \\ \log n & \leq & 2^i \\ \Updownarrow & & \\ \log \log n & \leq & i, \end{array}$$

giving the matching upper bound $O(\log \log n)$.

Now let us look at the general case with p processors. The algorithm tests for 3 cases:

- If $p \geq n^2$ the algorithm computes the maximum directly by comparing all pairs.
- If $n^2 > p \geq n$ the algorithm partitions the input elements into $\sqrt{p}$ groups and computes maximum of these groups (of size $\frac{n}{\sqrt{p}}$) in parallel by $\sqrt{p}$ recursive calls, each using $\sqrt{p}$ processors. In this way we can still compute the maximum of the return values in constant time.

To calculate the recursion depth let n_j be the number of elements per block (ignoring rounding) at depth h of the recursion. Then,

$$n_h \;:=\; \frac{n}{\Pi_{i=1}^{h} p^{2^{-i}}} \;=\; \frac{n}{p^{1-2^{-h}}}$$

The recursion returns when $n_h^2 \leq p^{2^{-h}}$ and there are enough processors to find the maximum in one step.

$$\begin{array}{rrcl}
 & \frac{n}{p^{1-2^{-h}}} & \leq & p^{2^{-h}} \\
\Updownarrow & & & \\
 & \log n & \leq & (1-2^{-j-1})(\log p) \\
\Updownarrow & & & \\
 & 2^{-j-1} & \leq & 1-\frac{\log n}{\log p} \\
\Updownarrow & & & \\
 & -j-1 & \leq & \log(1-\frac{\log n}{\log p}) \\
\Updownarrow & & & \\
 & \log(\frac{1}{1-\frac{\log n}{\log p}}) & \leq & j+1 \\
\Updownarrow & & & \\
 & \log(\frac{\log p}{\log(\frac{p}{n})}) & \leq & j+1 \\
\Uparrow & & & \\
 & 1+\log(\frac{\log n}{\log(\frac{p}{n})}) & \leq & j+1 \\
\Updownarrow & & & \\
 & \log(\frac{\log n}{\log(\frac{p}{n})}) & \leq & j
\end{array}$$

so we get the upper bound $\log(\frac{\log n}{\log(\frac{p}{n})})$ for the recursion depth, and thus the same upper bound $O(\log\log_{p/n} n)$ for the algorithm.

- If $p < n$, then the algorithm partitions the elements into p groups and computes the maximum sequentially in each group. It then uses p processors to find the maximum of the p remaining candidates in $O(\log\log p)$ time by the recursive algorithm described above. The total time is $O(\frac{n}{p} + \log\log p)$.

Solution 29: Without loss of generality assume that G is connected. By definition the faces of a planar graph G are 2-colourable if and only if the vertices of the dual graph G^* are 2-colourable, or in other words, G^* is bipartite.

If G^* is bipartite, all cycles in G^* must be of even length. Consider a vertex v and the circuit in G^* around it. This circuit must be even so the degree of v is also even. Since this is true for all vertices v we get that G Eulerian.

If G is Eulerian then it has no cut-edges and thus the faces of G^* are colourable. We argue that all cycles in G^* must be even and therefore G^* is bipartite. Consider an arbitrary cycle in G^*. We can assume without loss of generality that this cycle intersects each edge of G at most once since edges in G that intersect the cycle more than once can be subdivided by adding vertices between any two intersection points – what does not effect the graph being Eulerian.

In the embedding of G the cycle partitions the vertices of G to those inside the cycle and those outside. The length of the cycle is equal to the number of edges connecting vertices inside the cycle to vertices outside. But since all vertices have even degree, the number of edges between vertices that are inside the cycle and vertices that are outside is equal to the sum of the degrees of the vertices inside minus 2 times the number of edge with both ends inside – which is clearly an even number.

Solution 30 (Theis Rauhe): By the Four Colour Theorem we can partition the vertices of any planar graph into four independent sets. The two largest independent sets contains at least $|V|/2$ of G's vertices, implying the claim.

Part III

Linear Algebra in Combinatorics

Chapter 20

Invitation to Club Theory

20.1 A Tale of Two Cities

In a strange land there are two towns - Eventown and Oddtown. Both towns have 32 citizens who like to form lots of clubs. However, the rules for forming the clubs are a bit different in the two towns (reflected in their names).

Eventown:

1. Every club must consist of an **even** number of citizens.

2. Every two clubs must share an even number of members.

3. Two clubs can't have exactly the same members.

Oddtown:

1. Every club must consist of an **odd** number of citizens.

2. Every two clubs must share an even number of members.

(Observe that Oddtown doesn't need Eventown's rule 3 since this rule is a consequence of Oddtown rules 1 and 2.)

One may ask how many clubs can the citizens in the two towns form? In Eventown we can pair the 32 citizens into 16 couples and demand that no club break a couple. It is easy to see that all subsets of the 16 couples in the town are possible clubs and none of these clubs violate the rules. So Eventown can have up to $2^{16} = 65,536$ clubs[1]

For Oddtown we are going to show that the rules limit the number of clubs that can be formed quite drastically. We will use the most basic concepts from linear algebra, so let's start by briefly recalling these.

We are going to consider a vector space consisting of vectors $\mathbf{v} \in \mathsf{F}_2^k$, where F_2 is the finite field consisting of the two elements 0 and 1. Multiplication is as usual, and addition is modulo 2.

A set of vectors $\{\mathbf{v}_1 \cdots \mathbf{v}_m\}$ is *linearly dependent* if there exists a non–trivial linear relation:

$$\lambda_1\mathbf{v}_1 + \lambda_2\mathbf{v}_2 + \cdots \lambda_m\mathbf{v}_m = 0$$

[1] In fact Eventown can't have more than 65,536 clubs – see the exercises at the end.

(non–trivial means at least one $\lambda_i \neq 0$). A set of vectors is *linearly independent* if they are not linearly dependent. That is

$$\lambda_1 \mathbf{v}_1 + \lambda_2 \mathbf{v}_2 + \cdots \lambda_m \mathbf{v}_m = 0$$

implies that $\lambda_1 = \lambda_2 = \cdots = \lambda_m = 0$.

The *span* of $\{\mathbf{v}_1, \cdots, \mathbf{v}_m\}$ is the smallest subspace containing the vectors $\{\mathbf{v}_1, \cdots, \mathbf{v}_m\}$. Explicitly it can be written;

$$span(\mathbf{v}_1, \cdots, \mathbf{v}_m) = \{\lambda_1 \mathbf{v}_1 + \cdots + \lambda_m \mathbf{v}_m | \lambda_i \in \mathsf{F}_2\}$$

Note that for the vector space F_2^k we thus have $|span(\mathbf{v}_1, \cdots, \mathbf{v}_m)| \leq 2^m$. $B := \{\mathbf{v}_1, \cdots, \mathbf{v}_m\}$ is a *basis* for the vector space if the vectors in B are linearly independent and *span* $B = V$. One of the most basic facts in linear algebra is

Proposition 158 *Any two bases of V have the same cardinality, called the dimension of V.*

This invariant of a vector space is denoted $\dim_k V$ or simply dim V if the underlying field is clear from the context. A further basic fact is:

Proposition 159 *If $\mathbf{v}_1, \cdots, \mathbf{v}_k$ are linearly independent in a vector space of dimension n, then $k \leq n$.*

This might seem almost trivial, but we'll soon see how to extract very non–trivial mileage out of it!

Example 160 *$dim(\mathsf{F}_2^{32}) = 32$. Consider the 32 unit vectors:*

$$\begin{bmatrix} 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \cdots, \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix}$$

It is easy to see that they are linearly independent and span all of F_2^{32}.

Now we introduce an important additional structure in a vector space. In our space F_2^n, the *inner-product* (or *dot–product*) of two vectors $x = (x_1, \cdots, x_n)$ and $y = (y_1, \cdots, y_n)$ is $x \cdot y = x_1 y_1 + \cdots + x_n y_n$.

Consider m vectors $\mathbf{C}^1, \mathbf{C}^2, \cdots, \mathbf{C}^m$ in F_2^{32} where $\mathbf{C}^k$ is the bit–vector representation of club k in Oddtown: $\mathbf{C}_i^k = 1$ if citizen i is a member of club k and zero otherwise. In Oddtown the vectors $\mathbf{C}^i$ have the following properties. $\mathbf{C}^k \cdot \mathbf{C}^k = 1$ (in field F_2) since each club has an odd number of members. For $k \neq l$, $\mathbf{C}^k \cdot \mathbf{C}^l = 0$ since every two clubs share an even number of members. That is,

$$\mathbf{C}^i \cdot \mathbf{C}^j = [i = j]. \tag{20.1}$$

Now let us examine what happens if a linear relation

$$\lambda_1 \mathbf{C}^1 + \cdots + \lambda_m \mathbf{C}^m = 0,$$

holds. Taking the dot–product with $\mathbf{C}_i$ and using (20.1), we conclude that $\lambda_i = 0$. This holds for each i i.e. the vectors $\mathbf{C}^1, \cdots, \mathbf{C}^m$ are linearly independent implying that $m \leq 32$ by example 160 and fact 159. So the conclusion is that Oddtown can have at most 32 clubs!

Exercise 161 *Show that Oddtown can at least find a way to form 32 clubs!*

You have proved

Theorem 162 (Oddtown Theorem) *If Oddtown has n citizens, the largest number of clubs one can form is n.*

Exercise 163 *Go back and read the rules of Eventown. Show that Eventown cannot form more than 2^{16} clubs. In general, if Eventown has n citizens (for some positive integer n), then it is impossible to form more than $2^{\lfloor n/2 \rfloor}$ clubs. This is a bit difficult, so we shall develop the proof with hints along the way.*

1. *Recall from linear algebra that for a linear subspace U of a vector space V equipped with an inner product, the* **othogonal complement** *$U^\perp$ is the subspace $\{v \mid u \cdot v = 0 \text{ for all } u \in U\}$. For a subspace U of F_2^n, show that*

$$dim\ U + dim\ U^\perp = n.$$

2. *Let S be the set of all incidence vectors of clubs in Eventown, and let $U := span(S)$. Argue that the Eventown rules imply that U is a subspace of $U^\perp$.*

3. *Finally complete the proof by combining the two pieces above and using the observation we made that in F_2^n, the cardinality of a subspace of dimension k is exactly 2^k.*

You have proved:

Theorem 164 (Eventown Theorem) *If Eventown has n citizens, the maximum number of clubs that can be formed is $2^{\lfloor n/2 \rfloor}$.*

These mildly amusing examples served to illustrate the objects of our study for the next few lectures: *Set Systems*. In general, one has a structure $(S, \mathcal{F})$ where S is a finite set (we can canonically fix this set to be $[n]$ for some positive integer n) and $\mathcal{F}$ is a family of subsets satisfying some conditions. The object is to investigate how large this family can be under the given conditions. The subject is therefore called *Extremal Set Theory*, but we could, somewhat irreverently, call it *Club Theory*.

20.2 Phones, networks and addressing

Our focus in illustrating the use of linear algebra in combinatorics will be Club Theory, and we shall return to it shortly, but at this point, we want to show–off a bit with a diversion to hint at other ways in which linear algebra can be used in combinatorics.

In figure 20.1 there is an example of a telephone network. Here phone A wants to be connected to phone B. To achieve this, phone A must send a request through the other phones at the path to phone B. How can phone A figure out which direction to send its request without having any global knowledge of the topology of the network?

Suppose the network is a cube–like figure 20.2 with each vertex having coordinates $x = (x_1, x_2, \cdots, x_d) \in \{0, 1\}^d$. The Hamming distance between the two symbols in our alphabet,

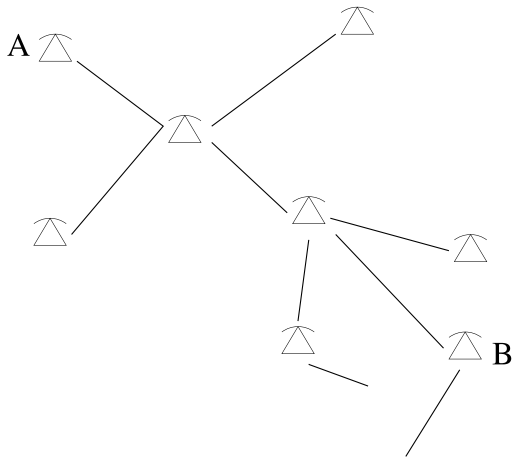

Figure 20.1: A telephone network

$d_H(a,b) := [a \neq b]$ extends to the Hamming distance between strings x and y over the alphabet naturally as follows:

$$d_H(x,y) = \sum_i d(x_i, y_i) = \sum_i [x_i \neq y_i].$$

An important property is that the Hamming distance between the two coordinates x and y equals the distance between their corresponding vertices in the cube - $d_H(x,y) = d_G(x,y)$. Suppose a phone with coordinate x is going to send a request to coordinate y. It may now send to one of its neighbours with a coordinate z such that $d_H(x,y) = d_H(z,y) + 1$. From z the distance is by the property decreased by one and continuing from z in the same way, the request reaches y through a shortest possible path. It is clear that this strategy only needs local information about neighbour coordinates and the address of the destination.

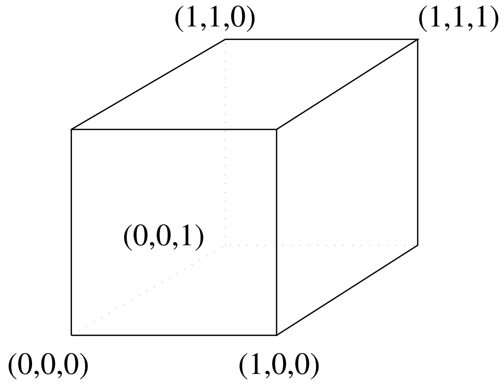

Figure 20.2: Cube

For general graphs $G = (V, E)$ the above method formalises into constructing a mapping:

$$f : V \to S^m, \quad S^m = \{(x_1, \cdots, x_m) \mid x_i \in \{0,1\}\}$$

such that $d_G(x,y) = d_H(f(x), f(y))$ for all x and y in V. Such a mapping is called an *addressing*.

Let's try out some special cases. For a tree $T = (V, E)$ there indeed exists an addressing $f : V \to S^{|V|-1}$. The proof is by induction on the number of vertices in the tree. The base case for trees with one or two vertices are trivial. Suppose we have addressing for trees upto size n, and let $T = (V, E)$ be a tree of size $n+1$. Let v be a leaf and let T' be the rest of T with v removed. We may now by the induction hypothesis give T' an addressing; $f' : V \setminus \{v\} \to S^{n-1}$. Let u be the vertex connected to v. Now define the addressing f for T to be:

$$f(v) = (1, y_1, y_2, \cdots, y_{n-1}), \text{ where } (y_1, \cdots, y_{n-1}) = f'(u)$$
$$f(w) = (0, y_1, y_2, \cdots, y_{n-1}), \text{ where } (y_1, \cdots, y_{n-1}) = f'(w)$$

It is easy to see that f is indeed the desired addressing.

It is not always possible to find an addressing for a graph. The simplest example of this is the triangle K_3. Suppose it had an addressing. Then each pair of vertices must have Hamming distance 1, which is impossible since $d_H(u, v) = 1$ and $d_H(v, w) = 1$ implies that $d_H(u, w) \in \{0, 2\}$.

To get around this problem we simply extend our addressing alphabet to contain a special joker symbol *, which matches both 0 and 1 (and itself). Formally for $a, b \in S_*$, $S_* := \{0, 1, *\}$:

$$d(a, b) = \begin{cases} 1 & \text{if } \{a, b\} = \{0, 1\} \\ 0 & \text{otherwise} \end{cases}$$

and the Hamming distance for coordinates S_*^m:

$$d_H(x, y) = \sum_i d(x_i, y_i)$$

In analogy to the previous a graph has an addressing for such extended coordinates if there exist a mapping $f : V \to S_*^m$ such that $d_G(x, y) = d_H(f(x), f(y))$.

Example 165 *The complete graphs K_n with $V = \{v_1, \ldots, v_n\}$ has the following addressing $f : V \to S_*^{n-1}$*

$$\begin{array}{rcl} f(v_1) & = & (1, *, *, *, \cdots, *) \\ f(v_2) & = & (0, 1, *, *, \cdots, *) \\ f(v_3) & = & (0, 0, 1, *, \cdots, *) \\ \vdots & \vdots & \vdots \\ f(v_{n-1}) & = & (0, 0, 0, \cdots, 0, 1) \\ f(v_n) & = & \underbrace{(0, 0, 0, \cdots, 0, 0)}_{n-1} \end{array}$$

Exercise 166 *Prove that it is possible to find an addressing $f : V \to S_*^m$ for any graph, provided that m is sufficiently large.*

The *squashed cube dimension* of G is the least m such that there is an addressing $f : V \to S_*^m$. We denote this by $d_*(G)$.

Our construction in the example above showed that $d_*(K_n) \le n - 1$. We will now prove that this is in fact optimal.

Before stating the proof we are going to introduce another concept from linear algebra. The *rank* of a matrix A is the dimension of the vector-space spanned by its row-vectors or equivalent the column-vectors. We will use the following *subadditivity* property of the rank: $rank(A+B) \leq rank(A) + rank(B)$.

A *tournament* $G := (V, E)$ is a directed graph with the property that for every pair of vertices u, v either $(u, v) \in E$ or $(v, u) \in E$ but not both. Let $D(G)$ denote the adjacency matrix for a graph G. The definition of a tournament means that its adjaceny matrix satisfies the equation:

$$A + A^T = J - I, \tag{20.2}$$

where J is the $(n \times n$)matrix of all 1s and I is the identity.

We claim that a matrix satisfying (20.2) has rank n or $n-1$. Suppose for contradiction that $rank(A) \leq n-2$. Then $rank \begin{bmatrix} A \\ J \end{bmatrix} \leq n-1$. So there must exist a $x \neq 0$ such that $Ax = 0$ and $Jx = 0$. But then we arrive at the following contradiction:

$$0 = x^T(A + A^T)x = x^T(Jx) - x^T I x = -x^T x < 0$$

Let $f : K_n \to S_*^m$ be an addressing of K_n with addresses of length m. Consider an edge (v, u) with $f(v) = (x_1, x_2, \cdots, x_m)$ and $f(u) = (y_1, y_2, \cdots, y_m)$. Since f is an addressing there exists exactly one index k where $f(u)$ and $f(v)$ differ. Each edge in E is labelled with such a k. Let H_k be the subgraph of edges labelled k. H_k is bipartite since its vertices can be divided into the two independent sets with vertices having a zero or one at index k respectively. Moreover, H_k is a complete bipartite graph and every two subgraphs $H_k, H_{k'}$ for $k \neq k'$ have disjoint edges. Thus we have a hence a decomposition $K_n = H_1 \cup H_2 \cup \cdots \cup H_m$ of K_n into edge–disjoint complete bipartite graphs. We are then left with determining the least m such that K_n can be covered by m complete bipartite graphs with disjoint edges.

Take such a decomposition $K_n = H_1 \cup H_2 \cup \cdots \cup H_m$. Let $\hat{H}_i$ be the digraph obtained from H_i by directing its edges from one the independent vertex sets to the other (arbitrarily). Note that $rank(D(\hat{H}_i)) = 1$ for each $i \in [m]$. Define the tournament T_n to be $\hat{H}_1 \cup \cdots \cup \hat{H}_m$. Since T_n is a tournament, $rank(D(T_n)) \leq n-1$, as observed a while back. Hence,

$$\begin{aligned} n-1 \quad \leq \quad rank(D(T_n)) &= rank(D(\hat{H}_1) + D(\hat{H}_2) + \cdots + D(\hat{H}_m)) \\ &\leq rank(D(\hat{H}_1)) + \cdots + rank(D(\hat{H}_m)) \\ &= 1 + \cdots + 1 = m. \end{aligned}$$

Thus our addressing for the complete graph was indeed optimal.

Chapter 21

Some Club Theory Classics

Today we would like to solve the following problem which might look at first sight to have been taken from a Calculus book: Maximise the integral

$$\int_0^{2\pi} \cos(m_1 x)\cos(m_2 x)\cdots\cos(m_n x)dx \tag{21.1}$$

where $n > 0$ is fixed and $\mathbf{m} = (m_0, m_1, \ldots m_n) \in \mathsf{N}^n$ can be chosen at will.

We can rewrite (21.1) by recalling the trigonometric formula

$$2\cos(x)\cos(y) = \cos(x+y) + \cos(x-y). \tag{21.2}$$

If you don't remember this, here's a refresher. Start with the magic formula of Euler:

$$e^{ix} = \cos x + i \sin x.$$

(For numerophiles, setting $x = \pi$ we get $e^{i\pi} = -1$, an equation relating some of the most important constants in Mathematics!) Thus,

$$\cos(x+y) = \mathrm{Re} e^{i(x+y)}, \quad \cos(x-y) = \mathrm{Re} e^{i(x-y)}.$$

Note that

$$e^{ix} + e^{-ix} = 2\cos(x).$$

Now we can calculate:

$$\begin{aligned}
\cos(x+y) + \cos(x-y) &= \mathrm{Re}(e^{i(x+y)}) + \mathrm{Re}(e^{i(x-y)}) \\
&= \mathrm{Re}(e^{i(x+y)} + e^{i(x-y)}) \\
&= \mathrm{Re}(e^{ix}(e^{iy} + e^{-iy})) \\
&= \mathrm{Re}(e^{ix} 2\cos(y)) \\
&= 2\cos(y)\mathrm{Re}(e^{ix}) \\
&= 2\cos(y)\cos(x)
\end{aligned}$$

This proves (21.2). We can put this in a more symmetric form by recalling that $\cos(-x) = \cos(x)$. Thus:

$$4\cos(x)\cos(y) = \cos(x+y) + \cos(x-y) + \cos(-x+y) + \cos(-x-y).$$

In general, one can easily show by induction that

$$2^n \cos(x_1)\cos(x_2)\ldots\cos(x_n) = \sum_{a_1,a_2,\ldots,a_n=\pm 1} \cos(a_1x_1 + a_2x_2 + \ldots + a_nx_n).$$

Now we can rewrite the formula under the integral from (21.1) in the following way:

$$\cos(m_1x)\cos(m_2x)\cdots\cos(m_nx) = \frac{1}{2^n} \sum_{(a_1,a_2,\ldots a_n)\in\{-1,1\}^n} \cos((a_1m_1 + a_2m_2 + \cdots + a_nm_n)x) \tag{21.3}$$

The integral (21.1) now becomes:

$$\begin{aligned} \int_0^{2\pi} \frac{1}{2^n} \sum_{a_1,\ldots,a_n=\pm 1} \cos((a_1m_1 + a_2m_2 + \cdots + a_nm_n)x)dx &= \\ \frac{1}{2^n} \sum_{a_1,\ldots,a_n=\pm 1} \int_0^{2\pi} \cos((a_1m_1 + a_2m_2 + \cdots + a_nm_n)x)dx & \end{aligned}$$

Now recall that for an integer p,

$$\int_0^{2\pi} \cos(px)dx = \begin{cases} 0, & \text{if } p \neq 0; \\ 2\pi, & \text{if } p = 0. \end{cases}.$$

Thus, our integral is

$$\begin{aligned} \frac{1}{2^n} \sum_{a_1,\ldots,a_n=\pm 1} \int_0^{2\pi} \cos((a_1m_1 + a_2m_2 + \cdots + a_nm_n)x)dx &= \\ \frac{2\pi}{2^n} \sum_{a_1,\ldots,a_n=\pm 1} [a_1m_1 + a_2m_2 + \cdots + a_nm_n = 0] & \\ &= \frac{2\pi}{2^n} N(n,\mathbf{m}) \end{aligned}$$

where $N(n,\mathbf{m})$ is the number of solutions to the equation

$$a_1m_1 + \ldots + a_nm_n = 0, \tag{21.4}$$

with $a_i = \pm 1, i \in [n]$. Thus we must pick $\mathbf{m} := (m_1, \ldots, m_n)$ to maximise the number of solutions to (21.4). Another way to look at (21.4) is that $N(n,\mathbf{m})$ is the number of subsets $I \subseteq [n]$ such that

$$(\mathbf{W}): \quad \sum_{i\in I} m_i = \sum_{i\in \bar{I}} m_i, \tag{21.5}$$

(where $\bar{I} := [n] \setminus I$). Indeed the correspondence between a solution $(a_1, \ldots, a_n)$ to (21.4) and a subset I satisfying (21.5) is that the former is the characteristic vector of the latter with the convention that

$$a_i = \begin{cases} +1 & \text{if } i \in I; \\ -1 & \text{if } i \notin I. \end{cases}$$

Focus then on the club system $([n], \mathcal{F})$ where each $I \in \mathcal{F}$ satisfies (W). Thus $N(n,\mathbf{m})$ the size of the largest such family $\mathcal{F}$. Set $N(n) := \max_{\mathbf{m}} N(n,\mathbf{m})$.

If $n = 2k$ is even , we can pick $m_1 = m_2 = \cdots = m_n = 1$ to get

$$N(2k) \geq \binom{2k}{k}. \tag{21.6}$$

If $n = 2k+1$ is odd, we can pick $m_1 = m_2 = \cdots = m_{n-1} = 1$ and $m_n = 2$

$$N(2k+1) \geq 2\binom{2k}{k-1}. \tag{21.7}$$

Is there anything better possible? We'll see shortly. Observe that the sets in our family $\mathcal{F}$ satisfy the following two properties:

(SYM) : $I \in \mathcal{F} \leftrightarrow \bar{I} \in \mathcal{F}$.

(A) : $I, J \in \mathcal{F}$ and $J \subseteq I \rightarrow I = J$.

The first property says that our family is symmetric and the second that the sets in our family are maximal. Let's throw away the weights for a moment and ask how large a family can we find to satisfy (SYM) and (A). Notice that an upper bound on such a family is automatically an upper bound on our original family.

Looking at the picture of the lattice of subsets of $[n]$, we see that our family $\mathcal{F}$ forms an *antichain*. How large an antichain can one find in this lattice? Notice that all the sets at any *level* in the lattice i.e. all subsets of a fixed size form an antichain. Taking all the subsets at the "middle" level, namely the sets of size $\lfloor n/2 \rfloor$, we see that we can form an antichain of size $\binom{n}{\lfloor \frac{n}{2} \rfloor}$. Can we get a larger antichain?

Let $\mathcal{F}_r$ denote the sets at level r in an antichain $\mathcal{F}$, for each $0 \leq r \leq n$. (Note that if $\mathcal{F}$ has at least two members then $\mathcal{F}_0$ and $\mathcal{F}_n$ are necessarily empty!). Let us take a level $r < \lfloor \frac{n}{2} \rfloor$ and "push it upwards". It is natural to consider the *upper shadow*:

$$\partial \mathcal{F}_r := \{ I \in \binom{[n]}{r+1} \mid J \subseteq I \text{for some } J \in \mathcal{F}_r . \}$$

Note that none of the sets in $\partial \mathcal{F}_r$ can be in $\mathcal{F}$. Suppose we exchange $\mathcal{F}_r$ for $\partial \mathcal{F}_r$. We do not violate the antichain property (A). Moreover,

$$|\partial \mathcal{F}_r| / |\mathcal{F}_r| \geq \frac{n-r}{r+1} \geq 1.$$

Hence the exchange only increases the size of our antichain. Similarly, for a level $r > \lfloor \frac{n}{2} \rfloor$, we can "push it downwards" without decreasing the size of our antichain. . Eventually we are left with an antichain of subsets all of size $\lfloor \frac{n}{2} \rfloor$ We have shown

Theorem 167 (Sperner's Lemma) *The largest antichain in the lattice of all subsets of a set of n elements has size $\binom{n}{\lfloor \frac{n}{2} \rfloor}$.*

Applied to our problem,we deduce that for $n = 2k$ even, $|\mathcal{F}| \leq \binom{2k}{k}$ matching the lower bound of (21.6).

Since this is a very beautiful classical result, we can't resist the temptation to give another proof. Consider the *maximal chains* in the lattice i.e. chains that start at $\emptyset$ and go all the way upto $[n]$. How many maximal chains are there? Clearly $n!$: starting with $\emptyset$ at the bottom, there are n choices for the next element, $n-1$ for the next and so on. Number these maximal chains as $C_1, C_2, \ldots, C_{n!}$. Suppose our antichain family $\mathcal{F}$ can be listed as $A_1, A_2, \ldots, A_m$. Form the $m \times n!$ matrix M:

$$M(i,j) := \begin{cases} 1, & \text{if } A_i \in C_j; \\ 0 & \text{otherwise.} \end{cases}.$$

Now look at a column (a chain). How many 1's are there in the column? At most one since our family is an antichain. Hence the sum of the entries in any column is at most 1 and so $\sum_{i,j} M(i,j) \leq n!$. What about a row sum? Take the element A_i corresponding to row i. How many chains can it be part of? The same argument as for counting the number of maximal antichains shows that ther are exactly $|A_i|!(n-|A_i|)!$ chains containing A_i. Hence the row sum for the ith row is $|A_i|!(n-|A_i|)!$. Summimg the matrix entries by row sums, we get $\sum_{i,j} M(i,j) = \sum_i |A_i|!(n-|A_i|)!$.Comparing the two different ways of forming the sum, we have $\sum_i |A_i|!(n-|A_i|)! \leq n!$ Dividing by $n!$ this reads

$$\sum_i 1/\binom{n}{|A_i|} \leq 1.$$

Now observing that $\binom{n}{|A_i|} \leq \binom{n}{\lfloor \frac{n}{2} \rfloor}$, we get Sperner's Lemma.

Now that we've settled the even case, let's fix attention on the odd case. For this, we may of course apply Sperner's Lemma. This yields an upper bound of

$$\binom{2k+1}{k} = \frac{2k+1}{k}\binom{2k}{k-1} > 2\binom{2k}{k-1},$$

which is just over our lower bound (21.7).

That should not be too surprising – after all, we dropped many conditions in going from (W) to (SYM). Consider the following *intersection* property:

(I): $I, J \in F \rightarrow I \cap J \neq \emptyset$.

A family satisfying (I) is said to be *intersecting.*

We claim: if $\mathcal{F}$ satisfies (S) then the condition (I) is equivalent to our original condition (A). Indeed,

$(A) \rightarrow (I)$: For $I, J \in \mathcal{F}$, $I \cap J = \emptyset. \leftrightarrow I \subseteq \bar{J}$. But by (SYM) $\bar{J} \in \mathcal{F}$ which contradicts (A).

$(I) \rightarrow (A)$: For $I, J \in \mathcal{F}$, $I \subseteq J \leftrightarrow I \cap \bar{J} = \emptyset$. But by (SYM) $\bar{J} \in \mathcal{F}$, contradicting (I).

For each $I, \bar{I} \in F$, keep the smaller to form a new club system $([n], \mathcal{F}')$ that is, $\mathcal{F}' := \{I \in \mathcal{F} \mid |I| \leq |\bar{I}|\}$,. Note that $|\mathcal{F}| = 2|\mathcal{F}'|$.

Now use the strategy of the first proof of Sperner's Lemma to get a family $\mathcal{G}$ which is at least as large as $\mathcal{F}'$ but all of whose elements belong to middle level $\binom{n}{\lfloor \frac{n}{2} \rfloor}$ of the lattice of subsets of $[n]$.

Theorem 168 (Erdos-Ko-Rado) *Let $([n], \mathcal{G})$ be a club system such that $\mathcal{G}$ is*

- *k–uniform i.e. each club in $\mathcal{G}$ has the same number k of elements.*

- *intersecting i.e. $\mathcal{G}$ satisfies (I).*

Then $|\mathcal{G}| \leq \binom{n-1}{k-1}$.

This is a true party theorem! A generous host invites the citizens to a party and seats his guests around a round table. The clubs make the special request that all the members of a club be placed consecutively. Of course not all such requests can be satisfied at any given seating. In fact an inspection of the seating on the table shows that due to the intersection conditions satisfied by our club system, at most k of our clubs get their wish fulfilled i.e. there are n clubs of k elements each seated consecutively around the table, but at most k of them belong to $\mathcal{G}$. The gallent host decides to be fair to one and all, and throws a series of $n!$ parties trying out a new seating each time. Of course after this exhausting regime, every club has its request satisfied – many times over! Moreover. it was a very equitable system: every club got its request fulfilled exactly the same number of times. However, at each seating, at most k of our clubs were satisfied. The only way to reconcile this is if the ratio $|\mathcal{G}|/\binom{n}{k}$ is at most k/n. That is,

$$|\mathcal{G}| \leq \frac{k}{n}\binom{n}{k} = \binom{n-1}{k-1}.$$

Thus for our family $\mathcal{F}$ satisfying (SYM) and (A), we have

$$|\mathcal{F}| = 2|\mathcal{F}'| \leq 2|\mathcal{G}| \leq 2\binom{2k}{k-1}.$$

Finally, let's not forget the problem we started with!

$$\max_{m_1,\dots,m_n} \int_0^{2\pi} \cos(m_1 x)\cos(m_2 x)\cdots\cos(m_n x)dx = \begin{cases} \frac{2\pi}{2^{2k}}\binom{2k}{k}, & \text{if } n = 2k; \\ \frac{2\pi}{2^{2k}}\binom{2k}{k-1}, & \text{if } n = 2k+1. \end{cases}.$$

In fact this is the very point at which the authors of the article on which this exposition was based goofed up, see

M. Katchalski, M.S. Klamkin, and A. Liu, "An Experience in Problem Solving", *American Mathematematical Monthly*, 88:6, pp. 551–556, 1981.

Chapter 22

More Club Theory

Last time we were in Eventown and Oddtown. Today we arrive in Equaltown: here the n citizens form clubs according to the rule: *any two clubs have the same number of members $\lambda \leq n$ in common.* How many clubs can we form?

Theorem 169 (Fisher's Inequality) *If $([n], \mathcal{F})$ is Equaltown, then $|\mathcal{F}| \leq n$.*

To prove this, we start by recalling some simple facts from linear algebra.. A matrix A has full rank iff the homogeneous system of linear equations $Ax = 0$ has only the trivial solution $x - 0$. A matrix A is called *positive definite* if :

$$x^T Ax > 0 \text{ for all } x \neq 0.$$

It is easily seen that a positive definite matrix has full rank, for a non–zero solution to $Ax = 0$ would imply that $x^T Ax = 0$ as well. A matrix A is called *positive semi-definite* if :

$$x^T Ax \geq 0 \text{ for all } x \neq 0.$$

Suppose C is a positive definite matrix, and D a positive semi-definite matrix. What about the matrix $C + D$? We compute for any $x \neq 0$,

$$x^T(C + D)x = x^T Cx + x^T Dx > 0, \tag{22.1}$$

since $x^T Cx > 0$ and $x^T Dx \geq 0$. Thus $C + D$ is positive definite as well.

We can now prove Fischer's inequality:

case 1 Some $C_i \in \mathcal{F}$ has size λ. This means, that the other clubs have exactly this club C_i as the common intersection of any two of them. It is easy to see that there can be at most $n - \lambda + 1$ clubs in this case.

case 2 Suppose there are $m > 0$ clubs and $|C_i| > \lambda$ for each $i \in [m]$. Define $\mu_i := |C_i| - \lambda$ (> 0) for each $i \in [m]$, and consider the $n \times m$ *incidence matrix* of the club system:

$$M(i, j) = \begin{cases} 1, & \text{if } i \in C_j; \\ 0, & \text{otherwise.} \end{cases}.$$

Consider the $m \times m$ matrix $A := M^T M$. The rules of Equaltown translate in matrix language into the relation

$$A = \lambda J + diag(\mu_1, \ldots, \mu_m),$$

where J is the $(m \times m)$ all 1's matrix. The claim is, that A has full rank $(=m)$. Consider first the matrix λJ :

$$\begin{aligned}
x^T(\lambda J)x &= \lambda x^T J x \\
&= \lambda x^T \begin{bmatrix} \sum x_i \\ \vdots \\ \sum x_i \end{bmatrix} \\
&= \lambda \left[x_1 \sum x_i + x_2 \sum x_i + \cdots + x_m \sum x_i\right] \\
&= \lambda \left(\sum x_i\right)^2 \\
&\geq 0.
\end{aligned}$$

This shows, that λJ is positive semi-definite. What about the other summand?

$$\begin{aligned}
x^T diag(\mu_1, \ldots, \mu_m)x &= x^T \begin{bmatrix} \mu_1 x_1 \\ \vdots \\ \mu_m x_m \end{bmatrix} \\
&= \mu_1 x_1^2 + \cdots + \mu_m x_m^2 \\
&> 0 \quad \text{for } x \neq 0.
\end{aligned}$$

The matrix $diag(\mu_1, \ldots, \mu_m)$ is thus positive definite. The observation (22.1) tells us, that the matrix A is positive definite and hence has full rank m. To complete the proof, we use the fact from linear algebra that rank is *sub–multiplicative*: $rank(AB) \leq \max(rank(A), rank(B))$. So

$$m = rankA \leq rankM \leq n.$$

We shall now give an application of Club Theory to a problem of Ramsey Theory which you encountered before in the graph theory part: does there exist an integer $r(k,l)$, such that any graph with $r(k,l)$ or more vertices has a k-clique *or* an l-independent set (or both). Equivalently, one acn phrase this as follows: does there exist an integer $r(k,l)$ such that for each $n \geq r(k,l)$, if we colour the edges of the complete graph K_n by two colours, red and blue, then there is either a complete red subgraph on r vertices (i.e. a set of k vertices, all connected to each other by red edges) or a complete blue subgraph on l vertices? Let's use club theory to give some answers.

We will consider the special case $r(t,t)$ for which we already have established the following bounds :

$$2^{t/2} \leq r(t,t) \leq 2^{2t-2}.$$

The lower bound (by Erdös) however is shown by a non-constructive proof, which although correct may seem a bit unsatisfying to a computer scientist. First we show the following beautifully simple construction:

Theorem 170 (Z. Nagy)

$$r(t,t) \geq \binom{t-1}{3}.$$

Proof: Consider the complete graph on $\binom{t-1}{3}$ vertices. Identify a node with a 3–element subset of $[t-1]$. Colour the edges in the graph by the following colouring rule χ:(see figure 15.3) : for $A, B \in \binom{[t-1]}{3}$:

$$\chi(A,B) = \begin{cases} \text{red} & \text{if } |A \cap B| = 1; \\ \text{blue} & \text{otherwise;} \end{cases}$$

Fisher's Inequality for Equaltown tells us, that this graph has no *red* K_t, and Oddtown tells us it has no *blue* K_t. □

However $\binom{t}{3}$ is far from Erdös' lower bound of $2^{t/2}$! In order to prove a tighter bound we need to visit even stranger towns! But first some notation: Let L be some set of integers, then:

$$n \in L \pmod p \quad \text{iff} \quad n = l \pmod p \text{ for some } l \in L.$$

We are now ready to enter *mod-p town*, $([n], \mathcal{F})$ where $\mathcal{F} = \{A_1, \ldots, A_m\}$ are the clubs subject to the following rules: For a prime p and a set of integers L,

1. $|A_i| \notin L \pmod p$.
2. $|A_i \cap A_j| \in L \pmod p$.

Oddtown can be seen as a special case of mod-p town, by choosing L as the set $\{0\}$ and $p = 2$.

Theorem 171 (Mod–p town Theorem) *If $([n], \mathcal{F})$ is mod-p town, then:*

$$|\mathcal{F}| \le \binom{n}{s} + \binom{n}{s-1} + \cdots + \binom{n}{0}.$$

where $s := |L|$.

Proof: Let the incidence vectors of the clubs be defined in the obvious way:

$$\mathbf{A_i} = (A_{i,1}, \ldots, A_{i,n}) \in \{0,1\}^n,$$

where $A_{i,j} = 1$ iff $j \in A_i$. Define the following (multi–variate) polynomials:

$$\begin{aligned} f_i(\mathbf{x}) &:= \prod_{l \in L} (\mathbf{x} \cdot \mathbf{A_i} - l) \\ &= \prod_{l \in L} (x_1 A_{i,1} + \cdots + x_n A_{i,n} - l). \end{aligned}$$

(where $\mathbf{x} := (x_1, \ldots, x_n)$. Applying the polynomials to the incidence vectors gives the following :

$$f_i(A_{j1}, \ldots, A_{jn}) = \begin{cases} 0 \pmod p & i \ne j \quad \text{(by rule 2 of mod-p town)} \\ \ne 0 \pmod p & i = j \quad \text{(by rule 1 of mod-p town)} \end{cases} \tag{22.2}$$

Now for another simple but very useful trick: since we are usually interested in 0/1 arguments in which case the boolean law $x^2 = x$ applies, we shall form the *multi-linear* versions of our polynomials, $\tilde{f}_i \leftrightarrow f_i$. As an example of this consider the polynomial $f =$

$x^4y^2z^7 + x^2z^3 - y + 5$. The new polynomial will then be $\tilde{f} = xyz + xz - y + 5$. Observe that (22.2) still holds for these new multi-linear polynomials:

$$\tilde{f}_i(A_{j1}, \ldots, A_{jn}) = \begin{cases} 0 \pmod p & i \neq j \quad \text{(by rule 2 of mod-p town)} \\ \neq 0 \pmod p & i = j \quad \text{(by rule 1 of mod-p town)} \end{cases} \tag{22.3}$$

Let $M(n, s)$ be the vector space of multi-linear polynomials in n variables of degree at most s over the finite field F_p. We now claim, that $\tilde{f}_i$, $i = 1 \ldots m$ are linearly independent in this space. To see this, consider a linear relation

$$\lambda_1 \tilde{f}_1 + \cdots + \lambda_m \tilde{f}_m = 0,$$

and evaluate it at $(A_{i1}, \ldots, A_{in})$. By equation (22.3) this is equal to :

$$\lambda_i \tilde{f}_i(A_{i1}, \ldots, A_{in}) = 0 \pmod p.$$

Again by equation (22.3) this means, that $\lambda_i = 0 \pmod p$, proving that the $\tilde{f}_i$'s are indeed linear independent.
All polynomials in $M(n, s)$ are on the form :

$$\tilde{f} = \sum_i \lambda_i x_{i_1} \cdots x_{i_r}.$$

This means that,

$$\{x_1^{\tau_1} x_2^{\tau_2} \cdots x_n^{\tau_n}\} \ , \ \tau_i \in \{0, 1\} \ , \ \tau_1 + \cdots + \tau_n \leq s.$$

form a basis of $M(n, s)$. Given this it is easy to see, that the dimention of $M(n, s)$ is :

$$\dim\ M(n, s) = \sum_{0 \leq k \leq s} \binom{n}{k}.$$

Hence the number of clubs $m \leq \sum_{0 \leq k \leq s} \binom{n}{k}$. □

We can now give a lower bound for $r(t, t)$ that is more respectable:

Theorem 172

$$r(t, t) > t^{\Omega(\log t / \log\log t)}.$$

Proof: Let p be the largest prime such that $2\binom{p^3}{p-1} < t$. Consider the complete graph K_N, where $N = \binom{p^3}{p^2-1}$. The vertices of this graph are the $(p^2 - 1)$ element subsets of $[p^3]$. Colour the edges of this graph by the following colouring rule χ:

$$\chi(A, B) = \begin{cases} \text{red} & \text{if } |A \cap B| \neq p - 1 \pmod p; \\ \text{blue} & \text{otherwise;} \end{cases}$$

The claim is, that this graph contains no monochromatic K_t. If it contained a *red* K_t the vertices $(E_i, i \in [t])$ would have the following characteristics :

$$\begin{aligned} |E_i \cap E_j| &\in \{0, \ldots, p-2\} \pmod p \\ |E_i| &= p^2 - 1 \equiv p - 1 \pmod p \\ |E_i| &\notin \{0, \ldots, p-2\} \pmod p. \end{aligned}$$

Choosing $\{0, \ldots, p-2\}$ as the set L, this is seen to be impossible by the mod-p town theorem, since $\sum_{k=0}^{s} \binom{N}{k} < 2\binom{N}{s}$ if $s < N/4$, which holds in our case, because $s = p-1 < p^{p^2-1} < N$ for $p \geq 2$.

In order to complete the second part of the proof, we need a corollary to the mod-p town theorem.

Corollary 173 *Let $([n], \mathcal{F})$ be a town. If $\mathcal{F}$ is uniform (i.e. all clubs have the same size) and there is a set L of s integers where :*

1. $|A_i| \notin L$.
2. $|A_i \cap A_j| \in L$.

then :

$$|\mathcal{F}| \leq \binom{n}{s} + \binom{n}{s-1} + \cdots + \binom{n}{0}.$$

This follows from mod-p town by choosing the set $\{0, \ldots, k-1\}$ as L and some prime $p > k$.

Assume now that the above graph contains a *blue* K_t ; then there are vertices $(E_i i \in [t])$ such that

$$\begin{aligned} |E_i \cap E_j| &\in \{p-1, 2p-1, \ldots, p^2-p-1\} \\ |E_i| &= p^2 - 1. \end{aligned}$$

Choosing $\{p-1, \ldots, p^2-p-1\}$ as the set L, this violates the corollary by the same arguments as above.

Finally, it remains to actually estimate N which is our lower bound. This was defined in a somewhat complicated way: it is $\binom{p^3}{p^2-1}$ where p is the largest prime such that $\binom{p^3}{p-1} < t$. We shall now estimate what this comes out to in terms of t i.e. $N = N(t)$. We shall do this in a somewhat lazy and inaccurate manner, leaving it to you to tidy up!

First, by the density of primes, there is always a prime between n and $2n$ for any positive integer n; so we may pretend that p is an integer rather than a prime. Next we will use the following elementary estimates of binomial coefficients:

$$(n/k)^k \leq \binom{n}{k} \leq (en/k)^k.$$

Thus $\binom{p^3}{p-1} \leq (ep^3/(p-1))^{p-1}$ and we may estimate p as the solution to the equation $(ep^2)^p = t$. In the exercise below, you are asked to show that this solution satisfies $p = \Omega(\ln t / \ln \ln t)$. Thus finally,

$$\begin{aligned} N &:= \binom{p^3}{p^2-1} \\ &\geq p^{p^2-1} \\ &= 2^{(p^2-1)\log p} \\ &= 2^{\Omega((\log t)^2/\log\log t)} \\ &= t^{\Omega(\log t/\log\log t)}. \end{aligned}$$

This is the promised super–polynomial lower bound for the Ramsey number $r(t, t)$. □

Exercise 174 *Show that the solution p^* to the equation $(ep^2)^p = t$ satisfies the following lower bound: $p^* = \Omega(\ln t / \ln\ln t)$. Proceed as follows: define two sequences $p_k, q_k, k \geq 0$ by the recurrences*

$$p_0 := 1, p_{k+1} := \ln t / (2 \ln q_k + 1),$$

$$q_0 := \ln t, q_{k+1} := \ln t / (2 \ln p_k + 1).$$

Show that

$$p_0 \leq p_1 \leq \cdots \leq p_k \leq \cdots \leq p^* \leq \cdots \leq q_k \leq \cdots \leq q_1 \leq q_0.$$

Finally argue that $p_k = \Omega(\ln t / \ln\ln t)$ and $q_k = O(\ln t / \ln\ln t)$ for each $k \geq 1$.

Chapter 23

Greedy Algorithms and Matroids

Let's start with a very familiar problem: Given a graph $G = (V, E, w)$ where w is a weight function $w : E \to \mathsf{R}_{\geq 0}$, find a spanning tree $T \subseteq E$ with minimum weight $w(T) := \sum_{e \in T} w(e)$.

This problem is an exception to the dictum that greed doesn't pay: We can use a *greedy algorithm* which colours the edges supposed to be in the tree *blue* and the other edges *red* using the following *greedy rules* repeatedly:

Red Pick a cycle without any red edges and colour the edge with maximum weight red.

Blue Pick a cutset with no blue edges and colour the edge with minimum weight blue.

Min Pick an uncoloured edge with minimum weight; if it does not form a cycle with the edges coloured blue so far, colour it blue otherwise colour it red.

Max Pick an uncloured edge of maximum weight; if it does not form a cut together with the red edges so far, then colour it red, otherwise blue.

We'll see that *any* sequence of applications of these rules results in a spanning tree of minimum weight! However, if our problem was instead to find a hamilton cycle of minimum weight, we would certainly not expect similar greedy rules to work! What's the difference between the two problems that accounts for this? When does greed pay and when does it not?

We shall turn to club theory to provide the answers to these questions. The rules we are about to list out define a special kind of club system $(S, \mathcal{F})$ called *matroids*. The reason why a greedy algorithm works for the spanning tree problem but not for the hamilton cycle problem is that in the former case, the underlying structure of the problem is a matroid (while in the latter it is not).

Definition 175 (Matroid: Independence Axioms) *$(S, \mathcal{F})$ is a* **matroid** *if*

I1: *$I \in \mathcal{F}$ and $J \subseteq I$ then $J \in \mathcal{F}$.*

I2: *$I_1, I_2 \in \mathcal{F}$ and $|I_2| = |I_1| + 1$ then there exist an $e \in I_2 \setminus I_1$ such that $I_1 \cup \{e\} \in \mathcal{F}$. (Exchange rule).*

The sets in the family $\mathcal{F}$ are called **independent** *sets.*

The inspiration for these rules comes from linear algebra: independent sets in a vector space satisfy exactly these conditions:

Example 176 (Vector spaces) *Let S be a vector space V and let $\mathcal{F}$ be a family of independent sets of vectors then $(S, \mathcal{F})$ is a matroid. Check out your linear algebra by verifying this!*

Example 177 (Matrices) *Let $M = [\mathbf{a}_1, \mathbf{a}_2, \cdots, \mathbf{a}_n]$ be a matrix (given by its column vectors). Put $S := \{\mathbf{a}_1, \mathbf{a}_2, \cdots, \mathbf{a}_n\}$ and let $\mathcal{F}$ be the subsets of columns that are linearly independent. Then $(S, \mathcal{F})$ is a matroid..*

Example 178 (Forests in graphs) *Let $G = (V, E)$, letting $S := E$ and $\mathcal{F}$:=$\{I \subseteq E \mid I$ is a forrest $\}$ then $(S, \mathcal{F})$ is a matroid.*

Exercise 179 (Node-disjoint paths in a graph (TSP)) *Let $G = (V, E)$, letting $S := E$ and $\mathcal{F} := \{p | p$ is a collection of node-disjoint paths $\}$. Is the structure $(S, \mathcal{F})$ a matroid? Axiom (I1) is certainly true, what about (I2)?*

Just as in vector spaces, an important notion is that of a *basis* or a *maximal independent set*.

Definition 180 (Maximal independent set) *Let $\mathcal{M} = (S, \mathcal{F})$ be a matroid then $I \in \mathcal{F}$ is a maximal independent set if I is set wise maximal in $\mathcal{F}$ i.e. $I \subseteq J \in \mathcal{F}$ implies that $I = J$.*

In terms of bases, the exchange rule (I2) from definition 175 can be reformulated as:

I2': For any subset $A \subseteq S$ any two maximal independent subsets of A have the same cardinality. That is for $I_1, I_2 \subseteq A$, if $I_1, I_2 \in \mathcal{F}$ then $|I_1| = |I_2|$.

Claim 181 *Given axiom (I1), the axioms (I2) and (I2') are equivalent.*

Proof: Let $(S, \mathcal{F})$ be a set system satisfying (I1).

I2 $\rightarrow$ I2': Let $I_1, I_2 \subseteq A$ be maximal independent. sets of a subset $A \subseteq S$ with $|I_1| \leq |I_2|$. If $|I_1| < |I_2|$, we may assume $|I_2| = |I_1| + 1$, otherwise use axiom (I1) to repeatedly delete elements from I_2 until this happens. Now from (I2) we can find an element $e \in I_2 \setminus I_1$ such that $I_1 \cup \{e\} \in \mathcal{F}$. However this contradicts maximality of I_1.

I2 $\leftarrow$ I2': Assume that the cardinality of any two maximal independence sets of any set A is the same. Now let $I_1, I_2 \in \mathcal{F}$ and $|I_2| = |I_1| + 1$. Define $A = I_1 \cup I_2$. Using (I2') we conclude that I_1 is not a maximal independent subset of A. Hence there is a $e \in I_2 \setminus I_1$ such that $I_1 \cup \{e\} \in \mathcal{F}$. □

In fact, one can alternatively characterise a matroid via its maximal independent sets:

Definition 182 (Matroid: Basis Axioms) *A set system $(S, \mathcal{B})$ is a matroid if the sets in $\mathcal{B}$ satisfies the following two conditions.*

B1: *For distinct $B_1, B_2 \in \mathcal{B}$ $B_1 \not\subseteq B_2$.*

B2: *Let $B_1, B_2 \in \mathcal{B}$ then for every $b_1 \in B_1$ there is a $b_2 \in B_2$ such that $B_1 \setminus \{b_1\} \cup \{b_2\} \in \mathcal{B}$*

Exercise 183 *Show that the two definitions are indeed equivalent:*

1. *Given a family $\mathcal{B}$ satisfying (B1) and (B2) show that there is a unique family $\mathcal{F}$ satisfying (I1) and (I2) such that $\mathcal{B}$ consists of exactly the maximal elements of $\mathcal{F}$.*

2. *Conversely given a family $\mathcal{F}$ satisfying (I1) and (I2), ther is a unique family $\mathcal{B}$ satisfying (B1) and (B2) compatible with $\mathcal{F}$.*

Axiom (I2') says that as in a vector space, any two bases have the same cardinality. In fact it is also possible to define a matroid from a *rank* function, where by $\text{rank}(A)$ we denote the (well–defined) size of a maximal independent subset of A.

Example 184 (Matroid: Rank Axioms) *The structure (S, r) where S a set, and $r : S \to \mathbb{N}$ a rank function is a matroid if the following rules are satisfied.*

R1 : $r(\emptyset) = 0$.

R2 : $A \subseteq S$ and $s \in S$ then $r(A) \leq r(A \cup \{s\}) \leq r(A) + 1$.

R3 : $r(A \cup B) + r(A \cap B) \leq r(A) + r(B)$.

Exercise 185 *Show that this definition is equivalent to any of the previous ones. (*HINT*: Given r, how would you define a set to be independent?)*

Let us now go back to our original problem of finding a spanning tree of minimum weight in a graph. From Example 178, we know that the underlying structure is a matroid. We shall therefore address the problem in the more general setup: let $\mathcal{M} := (S, \mathcal{F}, w)$ be a *weighted matroid* with a *weight function* $w : S \to \mathbb{R}$. We would like to find a maximal independent set $I \in \mathcal{F}$ of minimum weight $w(I) := \sum_{e \in I} w(e)$.

We shall formulate the more general versions of the greedy colouring rules in this setup. To do this, we need to introduce first the appropriate general versions corresponding to the concepts of cycles and cuts in graphs. A *cycle* $C \subseteq S$ in a matroid is a minimal dependent set i.e.

- $C \notin \mathcal{F}$.
- $C - \{x\} \in \mathcal{F}$ for every $x \in C$.

A *cut* $C' \subseteq S$ is a minimal subset that intersects all maximal independence sets i.e.

- $C' \cap I \neq \emptyset$ for every maximal $I \in \mathcal{F}$.
- $C' - \{x\} \cap I = \emptyset$ for every $x \in C'$ and some maximal $I \in \mathcal{F}$.

Now the general rules are exactly the same as before except that cycles and cuts are interpreted in this general sense.

There is a different way to define cuts that makes things more perspicuous. In fact it will also show that the Red/Blue and Min/Max rules are merely mirror images of each other, so it will half our work hereafter!

Definition 186 (Dual Matroid) *Let $\mathcal{M} := (S, \mathcal{B}, w)$ be a weighted matroid given via its basis elements. Define the* **dual** *matroid via its basis elements as follows: $\mathcal{M}^* := (S, \mathcal{B}^*, w^*)$ where $\mathcal{B}^* = \{B \subseteq S | S \setminus B \in \mathcal{B}\}$ and $w^*(e) = -w(e)$ for $e \in S$.*

Exercise 187 *Why did we define the dual via the basis elements? If $\mathcal{M} := (S, \mathcal{F})$ is a matroid, is the structure $\mathcal{M}' := (S, \mathcal{F}')$ a matroid, where $\mathcal{F}' := \{I \subseteq S \mid S \setminus I \in \mathcal{F}\}$? That is, is it true that a subset is independent in a matroid iff its complement is independent in the dual?*

Now observe the following *duality*:

- A cycle in $\mathcal{M}$ is a cut in $\mathcal{M}^*$ and vice-versa.
- The Blue/Red rule in $\mathcal{M}$ is the Red/Blue rule in $\mathcal{M}^*$ and the Min/Max rule in $\mathcal{M}$ is the Max/Min rule in $\mathcal{M}^*$.
- $I \subseteq S$ is a maximal independent set of minimal weight in $\mathcal{M}$ iff $S \setminus I$ is a maximal independent set of minimal weight in $\mathcal{M}^*$.

Applying our colouring rules gives us an *acceptable* partial colouring at any stage:

Definition 188 (Acceptable Colouring.) *An* **acceptable colouring** *is a pair (B, R) of subsets of S (corresponding to the elements coloured blue and red respectively) such that:*

- $B \cap R = \emptyset$.
- *B is independent in $\mathcal{M}$ i.e. $B \in \mathcal{F}$.*
- *R is independent in $\mathcal{M}^*$ i.e. $R \in \mathcal{F}^*$.*

An acceptable colouring is said to be **total** *if in addition,*

- $B \cup R = S$.

Exercise 189 *Let (B, R) be a total acceptable colouring. Show that B (R) is a maximal independent set in $\mathcal{M}$ $(\mathcal{M}^*)$.*

To talk about the weight we also have to define the weight of a colouring

Definition 190 (Weight of a colouring.) *Let (B, R) be an acceptable colouring, then the weight of the colouring is*

$$w(B, R) = \sum_{b \in B} w(b).$$

Now starting with the acceptable colouting $(B, R) := (\emptyset, \emptyset)$ we increase the two sets B and R using the greedy rules until we get a total acceptable colouring, hopefully of minimum weight.

Exercise 191 *Formulate the spanning tree with minimum weight problem in this language.*

We now come to the heart of the matter:

Lemma 192 (Extension Lemma) *If an acceptable colouring (B, R) has a total acceptable colouring $(\hat{B}, \hat{R})$ of min. weight, then any single application of the four greedy rules to (B, R) gives us an acceptable colouring (B', R') which also has a total acceptable colouring $(\tilde{B}, \tilde{R})$ of min. weight.*

Proof: What we want to show is the correctness of the right hand part in the following diagram

$$\begin{array}{ccc} (B, R) & \hookrightarrow & (\hat{B}, \hat{R}) \\ rule \downarrow & & \|w \\ (B', R') & \hookrightarrow & (\tilde{B}, \tilde{R}) \end{array}$$

where $\hookrightarrow$ denotes an extension of a colouring to a total acceptable colouring.

Let us show it for the rules Red and Min, this is enough because of duality.

Red Suppose that one application of the red rule gives an acceptable colouring $(B', R') = (B, R \cup \{e\})$. We would like to extend (B', R') to a total acceptable colouring of minimum weight. If $e \in \hat{R}$, then we are done, for $(\hat{B}, \hat{R})$ is the desired extension. If $e \in \hat{B}$, we would like somehow to *exchange* it with an element from $\hat{R} \setminus R$. As $(\hat{B}, \hat{R})$ are maximal independent sets $\hat{R} \cup \{e\}$ is dependent and then there exists a cycle $C \subseteq \hat{R} \cup \{e\}$. In that cycle pick an edge $e' \in C \setminus R$; this is always possible as R was independent. If we exchange e and e', we hope the result is a total acceptable colouring as well. The Exchange Lemma 194 confirms this. The result, $(\tilde{B}, \tilde{R}) = (\hat{B} \setminus \{e\} \cup \{e'\}, \hat{R} \setminus \{e'\} \cup \{e\})$ has weight

$$w(\tilde{B}, \tilde{R}) = w(\hat{B}, \hat{R}) - w(e) + w(e')$$

Because e had maximun weight $w(e) - w(e') \leq 0$ and this extension of (B', R') is a total acceptable colouring of min. weight.

Min Suppose one application of the Min rule gives an acceptable colouring $(B', R') = (B \cup \{e\}, R)$. As before, we would like to extend this to a total acceptable colouring of minimum weight. If $e \in \hat{B}$, we are already home. Else $e \in \hat{R}$. As before we get that $\hat{B} \cup \{e\}$ is dependent. Then there exists a minimal dependent set $C \subseteq \hat{B} \cup \{e\}$. Let $e' \in C \setminus B \cup \{e\}$ then we can exchange e, e' and still have a total acceptable colouring by the Exchange Lemma. The weight of this extension is $w(\hat{B}, \hat{R}) + w(e) - w(e') \leq w(\hat{B}, \hat{R})$ as e was picked to be of minimum weight.

□

Exercise 193 *Show that the cycle referred to in the proof above is in fact unique i.e. if I is a maximal independent set and $e \notin I$, then there is a unique cycle in $I \cup \{e\}$ containing e. This is called the* **fundamental cycle** *of I and e.*

Lemma 194 (Exchange Lemma) *Let $(\hat{B}, \hat{R})$ be a total acceptable colouring, $x \in \hat{R}$ and $C \subseteq \hat{B} \cup \{x\}$ a minimal dependent set. Then for all $y \in C$, $(\hat{B} \setminus \{y\} \cup \{x\}, \hat{R} \setminus \{x\} \cup \{y\})$ is total acceptable.*

Proof: If $x = y$ there is nothing to prove. Otherwise take $C \setminus \{y\}$ which is independent and note that $x \in C \setminus \{y\}$. Now keep adding elements from $\hat{B}$ until you have a maximal independent set. That set must in fact be $\hat{B} \setminus \{y\} \cup \{x\}$, because any two maximal independent sets have the same size. □

Starting with the acceptable colouring $(\emptyset, \emptyset)$ which certainly has a total acceptable extension of minimum weight, the Extension lemma guarentees that each application of the greedy rules yields a larger partial acceptable colouring that also has a total acceptable colouring of minimum weight. Eventually therefore, we in fact arrive at a total acceptable colouring of minimum weight. .

Theorem 195 *Let $\mathcal{M} := (S, \mathcal{F}, w)$ be a weighted matroid. Then applying a sequence of our greedy rules in arbitrary order results in a maximal independent set of minimum weight.*

In particular, our original greedy rules find a spanning tree of minimum weight in a weighted graph.

Exercise 196 *Note that at any stage the Min/Max rules are always applicable. Show that at any stage either the Red or the Blue rule is also applicable.*

There is a revealing converse to this: if the greedy rule works, then there must be an underlying matroid structure!

Proposition 197 *Let $\mathcal{M} := (S, \mathcal{F})$ be a set system satisfying (I1). If the greedy rules always find a maximal independent set of minimum weight for any weight function $w : S \to \mathsf{R}$, then $\mathcal{M}$ is in fact a matroid.*

Proof: If not, then there are independent sets I_1, I_2 violating (I2). By using (I1), fix such a pair I_1, I_2 with $|I_1| = |I_2| - 1 =: p$. Now assign weights as follows:

$$w(e) := \begin{cases} -(p+2) & \text{if } e \in I_1 \ ; \\ -(p+1) & \text{if } e \in I_2 \setminus I_1; \\ 0 & \text{otherwise;} \end{cases} .$$

Let us now apply our greedy algorithm and see what happens. The Min rule can be applied to pick first all the elements in I_1. Now however by our assumption, it can only pick elements outside of I_2. Thus the final maximal independent set produced has weight $-p(p+2) > -(p+1)^2 = w(I_2)$. But this contradicts our assumption that any sequence of applications of the greedy rules produce a maximal independent set of minimum weight. □

Exercise 198 *In the proof above, we used negative weights, which might be considered a bit artificial. Given a proof with non–negative weights.*

Chapter 24

Probability Spaces with Limited Independence

Let $X_1, \ldots, X_n$ be random variables over a sample space $(\Omega, \Pr)$. These are *independent* if for any $I \subseteq [n]$ and any reals $a_i, i \in I$,

$$Pr[\bigwedge_{i \in I} X_i = a_i] = \prod_{i \in I} Pr[X_i = a_i].$$

Independence of random variables is a very nice property and is used crucially in the proofs of many central theorems in classical probability theory, for instance look at the classical book by William Feller, *An Introduction to Probability Theory and its Applications*.

However, in Computer Science, we know that true random bits are very rare to come by. Thus we can view randomness as a *resource*, available to an algorithm. We would like to get away using as little randomness as possible, without losing too much in efficiency – ideally with no randomness at all, in which case we would have a deterministic algorithm. In fact, when we analyse various randomised algorithms carefully, we see that the random variables involved in the analysis are not really required to have the full independence property stated above, but instead somewhat weaker properties often suffice and this enables us to decrease the amount of randomness that our algorithms can get by with.

One such property is k-wise independence: The random variables $X_1, \ldots, X_n$ are *k-wise independent* if for any $I \subseteq [n], |I| \leq k$ and any reals $a_i, i \in I$,

$$Pr[\bigwedge_{i \in I} X_i = a_i] = \prod_{i \in I} Pr[X_i = a_i].$$

The word "independence" might suggest a connection to the familiar notion of vector space independence:

Proposition 199 *Let $\mathbf{v}_1, \ldots, \mathbf{v}_n \in \{0,1\}^d = \mathsf{F}_2^d$ such that any k of them are linearly independent. Consider the sample space (Ω, Pr) with $\Omega := \{0,1\}^d$, and $\Pr(\omega) = 1/2^d$ for each $\omega \in \Omega$. Define the following random variables :*

$$\begin{array}{rcl} X_i : \Omega & \to & \{0,1\} \\ \omega & \mapsto & \mathbf{v}_i \cdot \omega \end{array}$$

Then :

1. *For each* $i \in [n]$, $Pr[X_i = 0] = 1/2 = Pr[X_i = 1]$.
2. $X_1, \ldots, X_n$ *are* k*-wise independent.*

Let us look at the first assertion which asserts that the random variables we construct are in fact unbiased. Why is this? Fix some $i \in [n]$; by definition,

$$Pr[X_i = 0] = |\{\omega : \mathbf{v}_i \cdot \omega = 0\}|/2^d =: |\Omega_0|/2^d,$$

and

$$Pr[X_i = 1] = |\{\omega : \mathbf{v}_i \cdot \omega = 1\}|/2^d =: |\Omega_1|/2^d.$$

Now, there here exists a bijective mapping between Ω_0 and Ω_1 as follows: Each $\mathbf{v}_i$ is a nonzero vector and therefore it must contain at least one 1. Define the mapping $' : \Omega \to \Omega$ which flips the bit in the position where $\mathbf{v}_i$ is 1. Then :

$$\mathbf{v}_i \cdot \omega = 0 \quad \Leftrightarrow \quad \mathbf{v}_i \cdot \omega' = 1.$$

It is easy to see that $'$ is an injective mapping between Ω_0 and Ω_1 which proves the first statement.

Next we must prove that the random variables are k–wise independent, i.e. for any $I \subseteq [n], |I| \leq k$ and any reals $a_i, i \in I$,

$$Pr[\bigwedge_{i \in I} X_i = a_i] = \prod_{i \in I} \Pr[X_i = a_i],$$

i.e., from what we just proved,

$$Pr[\bigwedge_{i \in I} X_i = a_i] = 2^{-|I|}.$$

What do we have here? A system of linear equations in ω:

$$\mathbf{v}_i \cdot \omega = a_i, \quad i \in I.$$

We can write this in matrix language as:

$$A \cdot \omega = \mathbf{a}, \tag{24.1}$$

where $A = (\mathbf{v}_i, i \in I)$ and $\mathbf{a} = (a_i, i \in I)$. What can we say about the number of solutions to this linear system?

In order to reason about this, we need a few facts from linear algebra. Recall that the *kernel* of a matrix A is

$$\ker\, A := \{\mathbf{x} : A\mathbf{x} = 0\}.$$

If $A\mathbf{x}_1 = \mathbf{a}$ and $A\mathbf{x}_2 = \mathbf{a}$ then $A(\mathbf{x}_1 - \mathbf{x}_2) = 0$ i.e. $\mathbf{x}_1 - \mathbf{x}_2 \in \ker\, A$. Conversely if $A\mathbf{x} = \mathbf{a}$ and $A\mathbf{x}_h = 0$ (i.e. $\mathbf{x}_h \in \ker\, A$) then $A(\mathbf{x} + \mathbf{x}_h) = \mathbf{a}$.
Thus all the solutions to the system (24.1) can be described as follows: let $\mathbf{x}_{nh}$ be any fixed solution, then all other solutions are given by $\mathbf{x}_{nh} + \mathbf{x}_h$ for all distinct solutions $\mathbf{x}_h$ of the system $A\mathbf{x} = 0$. (The reason for the subscripts is that the equation $A\mathbf{x} = 0$ is the *homogeneous* counterpart of the *non–homogeneous* system (24.1). Thus there is a one to one correspondence between the solutions of (24.1) and elements of the kernel.

Thus we may concentrate on finding the cardinality of ker A. The nice thing about this is that ker A is a subspace – check! What is its dimension? By linear algebra this is simply $n - r$ where $r :=$ rank A. Finally, $|\ker A| = 2^{d-r}$.

What is the rank A for the system (24.1)? This is where the assumption of k–wise independence comes in: since any k of our vectors are linearly independent, the matrix A has full rank, $|I|$. Hence the number of solutions to (24.1) is $2^{n-|I|}$ and our probability is $2^{d-|I|}/2^d = 2^{-|I|}$ as desired.

But how can we get our hands on a set of n vectors such that any k of them are linearly independent? We shall show a simple construction for $k = 3$ although this can be generalised for any fixed k as well.

Let $\mathbf{v}_1, \ldots, \mathbf{v}_n$ be vectors in F_2^d given by $\mathbf{v}_i := (i_1, \ldots, i_d, 1)$ where $(i_1, \ldots, i_d)$ is the binary representation of i for $i \in [n]$ and $d = \lceil \log n \rceil$.

The claim is that these vectors are 3-wise linear independant :

1. Since no vector is zero, they are 1-independent.

2. Consider a relation $\lambda_i \mathbf{v}_i + \lambda_j \mathbf{v}_j = 0$ for $i \neq j$. Neither coefficient λ_i can be zero else we are in the first case. Thus $\mathbf{v}_i + \mathbf{v}_j = 0$, which is not true since $i \neq j$ and thus there is at least one position where the binary representations of i and j differ. Thus the vectors are 2–wise independent.

3. Finally let us see if a relation $\lambda_i \mathbf{v}_i + \lambda_j \mathbf{v}_j + \lambda_k \mathbf{v}_k = 0$ an hold for distinct i, j, k. No coefficient can be zero by the previous case. Thus we must have $\mathbf{v}_i + \mathbf{v}_j + \mathbf{v}_k = 0$ But this is impossible because of the 1 in the last position! Thus the vectors are 3–wise independent.

What is the size of the sample space for this construction? It is $2^{\lceil \log n \rceil + 1} \leq 4n$. Contrast this with the case of full independence where the sample space has size 2^n. In fact the construction we showed was used to *derandomise* a class of randomised algorithms.

We shall now show that our construction is essentially optimal.

Proposition 200 *Let d be a fixed positive integer and let $Y_1, \ldots, Y_n$ be d-wise independent random variables all with variance greater than zero (such variables are said to be not* almost constant*) on a sample space (Ω, Pr).*
Then $|\Omega| \geq c \cdot n^{(\lfloor d/2 \rfloor)}$ for some constant c.

Proof: Without loss of generality assume that $E[Y_i] = 0$ for all i (else consider $Y_i - E[Y_i]$). Since $var(Y_i) > 0$ this also means that $E[Y_i^2] > 0$. For each subset $S \subseteq [n]$, define the random variables :

$$Z_S = \prod_{i \in S} Y_i.$$

Consider the variables Z_S corresponding to sets S of size at most d:

$$\begin{aligned} E[Z_S^2] &= E[\prod_{i \in S} Y_i^2] \\ &= \prod_{i \in S} E[Y_i^2] \\ &> 0. \end{aligned} \tag{24.2}$$

In the second line we use the fact that the variables Y_i are d–wise independent. Similarly for subsets $S \neq T$, $\quad |S \cup T| \leq d$, we have,

$$\begin{aligned} E[Z_S Z_T] &= E[\prod_{i \in S\cap T} Y_i^2 \prod_{i \in S\cup T \setminus S\cap T} Y_i] \\ &= \prod_{i \in S\cap T} E[Y_i^2] \prod_{i \in S\cup T \setminus S\cap T} E[Y_i] \\ &= 0. \end{aligned} \tag{24.3}$$

The Z_S can be regarded as elements of the vector space R^Ω of real–valued functions on Ω. Consider the family of subsets $S_1, S_2, \ldots S_m \subseteq [n]$ such that for all $i \neq j$, $|S_i \cup S_j| \leq d$. Then we claim that $Z_{S_1}, \ldots, Z_{S_m}$ are linearly independent. For suppose we have a linear relation

$$\lambda_1 Z_{S_1} + \cdots + \lambda_m Z_{S_m} = 0.$$

Multiply by Z_{S_i} and take expectations using (24.2) and (24.3) to conclude that $\lambda_i = 0$ for each $i \in [n]$.

We conclude that

$$m \leq \dim\ \mathsf{R}^\Omega = |\Omega|.$$

Thus we are left with the following combinatorial problem of maximising the number of sets m which is left for you as an exercise:

Exercise 201 *Let $(S, \mathcal{F})$ be a club system with the rule that for any two $S, T \in \mathcal{F}$, we have $|S \cup T| \leq d$ for some fixed positive integer d. Then show that we can get a system with $|\mathcal{F}| = \Omega(n^{\lfloor d/2 \rfloor})$. Can you extend it to non–constant d? What about $d = n$?*

□

Recent Publications in the BRICS Lecture Series

LS-95-4 **Dany Breslauer and Devdatt P. Dubhashi.** ***Combinatorics for Computer Scientists*****. August 1995. viii+184 pp.**

LS-95-3 **Michael I. Schwartzbach.** ***Polymorphic Type Inference*****. June 1995. viii+24 pp.**

LS-95-2 **Sven Skyum.** ***Introduction to Parallel Algorithms*****. June 1995. viii+16 pp.**

LS-95-1 **Jaap van Oosten.** ***Basic Category Theory*****. January 1995. vi+75 pp.**