

Preface

One technique for proving the existence of an object with certain properties is to show that a random object chosen from an appropriate probability distribution has the desired properties with positive probability. This approach is known as the *probabilistic method*; using it to find graph colourings with special properties is the topic of this book.

The probabilistic method was pioneered and championed by Paul Erdős who applied it mainly to problems in combinatorics and number theory from 1947 onwards. The authors have been told that at every combinatorics conference attended by Erdős in the 1960s and 1970s, there was at least one talk which concluded with Erdős informing the speaker that almost every graph was a counter-example to his conjecture. Although this story is apocryphal, it does illustrate three facts about the probabilistic method which are worth bearing in mind.

The first is that the probabilistic method allows us to consider graphs which are both *large* and *unstructured*. Even a modern computer looking for a counterexample via exhaustive search would not be able to run through all the graphs on twenty nodes in a reasonable amount of time. A researcher in the 60s certainly cannot be expected to have done so. In contrast, the counterexamples constructed using the probabilistic method routinely contain many, say 10^{10} , nodes.

Furthermore, any attempt to build large counterexamples via an explicit construction necessarily introduces some structuredness to the class of graphs built, which thus restricts the graphs considered. To illustrate this, we remark that even though the clique and stability numbers of a typical graph on n vertices are both $O(\log n)$, we do not know how to efficiently construct a graph with clique and stability numbers which are this small and for many years we could do no better than $o(\sqrt{n})$. Thus, the probabilistic method allows us to boldly go where no deterministic method has gone before, in our search for an object with the properties we desire. This accounts for its power.

The second moral of our story is that the method is not just powerful, it is also easy to use. Erdős would routinely perform the necessary calculations to disprove a conjecture in his head during a fifteen minute talk. Indeed, in the introduction to his Ph.D. thesis (which was published in 1970), Vasek Chvátal

wrote in this thesis we study hypergraphs using the probabilistic method where by the probabilistic method we mean the taking of sums in two different ways. This *First Moment* approach was one of the main techniques used in the early days. For example, to bound the expected number of cliques of size k in a typical graph on n vertices, we consider each set S of k vertices and ask: in what proportion of the graphs is S a clique? Since this is roughly $2^{-\binom{k}{2}}$ (as each of the $\binom{k}{2}$ edges will be present with probability $\frac{1}{2}$), there are on average $\binom{n}{k}2^{-\binom{k}{2}}$ cliques of size k in a graph on n vertices. This manner of calculating the sum is much easier than the alternative approach, which consists of determining the proportion of graphs with no cliques of size k , those with one clique of size k , etcetera. Furthermore, this one line calculation implies that the clique number (and in the same vein stability number) of a typical graph is $O(\log n)$, while, as mentioned above, we do not know how to prove such graphs exist without recourse to the probabilistic method.

The third and final moral of our story about Erdős and his counterexamples is that in classical applications of the probabilistic method, the result obtained was not just that a positive proportion of the random objects had the desired property, but rather that the overwhelming proportion of these objects did. E.g. Erdős did not say some graph is a counterexample to your conjecture, but rather *almost every* graph is a counter-example to your conjecture. It would be fair to say that the classical technique was typically limited to situations in which the probability that a randomly chosen object had the desired property was fairly large.

There were two motivating factors behind our decision to write this book. The first was to provide a gentle introduction to the probabilistic method, so that researchers who so desired could add this powerful and easy to use weapon to their arsenal. Thus, the book assumes little or no background in probability theory. Further, it contains an introductory chapter on probability, and each of the tools we use is introduced in a separate chapter containing examples and (at least in the early chapters) exercises illustrating how it is applied.

Of course, new probabilistic tools have been introduced in the last half-century which are more sophisticated than switching the order of two summation signs. Many of these focus on the typical difference between a random non-negative real-valued variable X and its average (or *expected*) value, denoted $\mathbf{E}(X)$. Since, $\mathbf{E}(X)$ is often easy to compute using the first moment method, any observations we can make about $|X - \mathbf{E}(X)|$ translate into statements about the value of X . Results which bound $|X - \mathbf{E}(X)|$ are called *concentration* results because they imply that independently sampled values of X tend to be concentrated around $\mathbf{E}(X)$ (the reader may find the image of darts concentrated around the bulls-eye of a target helpful here).

One classical approach to concentration is the Second Moment Method which bounds the expected value of $|X - \mathbf{E}(X)|$ using the First Moment Method applied to a related variable. It is this approach which allows us to

tie down precisely the behaviour of the clique size of a typical graph on n vertices.

Certain random variables are better-behaved than others and for such random variables we can obtain much more information about $|X - \mathbf{E}(X)|$ than the relatively weak bounds on its expected value we obtain using the Second Moment Method. For example if Y is the number of heads obtained in n tosses of a *fair* coin then the probability that Y is k is simply $\binom{n}{k}2^{-n}$. Using this fact, and manipulating the terms involved, Chernoff obtained very strong bounds on the probability that $|Y - \mathbf{E}(Y)|$ exceeds t for a given real t . His bounds also apply to the number Y of heads obtained in n flips of a coin which comes up heads with some probability p .

Note that changing the outcome of any one coin flip can affect this number Y of heads by at most one. Recently, researchers have come up with more general tools which bound the concentration of random variables X which are determined by a sequence of n events such that changing the outcome of one event can affect the (expected) value of X by at most one. We present one such tool developed by Azuma, and another developed by Talagrand. Although the proofs that these bounds hold are non-trivial, as the examples given in the book attest, applying them is straightforward.

As mentioned above, using the classical probabilistic method researchers typically proved that a random object had certain desirable properties with probability near one. The Lovász Local Lemma is a very important tool which allows us to prove the existence of objects with certain properties when the probability that a random object has the desired properties is exponentially small. The proof of this lemma is short and it is very easy to use, as our examples once again attest. However, in contrast with e.g. the First Moment Method, it is difficult to construct efficient algorithms to find the objects the lemma guarantees exist. This is not too surprising as the number of such objects is so small. (In contrast, if almost every object has the property then to find such an object we can just pick one at random.) The last two chapters of the book discuss how to construct such algorithms.

Our second motivation for writing the book was to provide a unified treatment of a number of important results in graph colouring which have been proven using an iterative approach (the so-called semi-random method). These include results of Kim and Johansson on colouring triangle-free graphs with $O(\log n)$ colours, Kahn's results on $(1 + o(1))\Delta$ edge list colouring linear hypergraphs, Kahn's proof that asymptotically the Goldberg-Seymour conjecture holds, and the fact that there is an absolute constant C such that every graph of maximum degree Δ has a $\Delta + C$ total colouring.

Our treatment of the results discussed above shows that each of them can be proven via an application of the semi-random method where we analyze an iteration using the Local Lemma and our concentration results. Although the proofs all have a similar flavour (indeed we reworked all the proofs to illustrate the unity of this body of work and hopefully to make them more accessible),

some of them are more complicated than others, requiring auxiliary notions (such as entropy and hardcore distributions) and quite involved arguments. To ease the reader's burden, we present the easier results first, building up to more and more sophisticated variants of the technique. Fittingly, we end with Kahn's result on the list chromatic index of multigraphs which really is a tour-de-force.

Three researchers who have had a very important impact on the mathematical careers of the authors are Vasek Chvátal, Alan Frieze, and Colin McDiarmid. We both met all three of these researchers as graduate students. Vasek was the second author's Ph.D. supervisor and Alan supervised the first author. The first paper written by the second author on probabilistic combinatorics was joint with Alan and Colin. We are immensely grateful for all that we have learnt from these three researchers and the many enjoyable hours we have spent in their company.

We are very grateful to those researchers who read portions of the book, pointing out errors and suggesting reworkings. The following is a list of some of those researchers, we apologize to any we may have inadvertently left out.

Dimitris Achlioptas, Noga Alon, Spyros Angelopoulos, Etienne Birméle, Adrian Bondy, Babak Farzad, Nick Fountoulakis, Alan Frieze, Paul Gries, Jeff Kahn, Mark Kayll, Jeong Han Kim, Michael Krivelevich, Andre Kündgen, Mohammad Mahdian, Colin McDiarmid, Steve Myers, Ioannis Papoutsakis, Ljubomir Perkovic, Mohammad Salavatipour, Benny Sudakov, Frank Van Bussel, Van Vu.

The authors are very grateful to Cynthia Pinheiros Santiago who prepared the figures and converted the manuscript into the Springer Latex format.

The second author would also like to thank the University of Toronto, McGill University, and the University of Sao Paulo, which all provided generous hospitality and support during the writing of the book. Both authors would like to thank the Federal University of Ceara in Fortaleza, Brazil where much of the polishing of this book took place.

Despite the many careful readings and rereadings, errors will inevitably remain, however we feel that the manuscript is now ready for publication and Boy! are we glad to be finished.

Montreal
September 2001

Michael Molloy
Bruce Reed

10. Talagrand's Inequality and Colouring Sparse Graphs

10.1 Talagrand's Inequality

In Chap. 5 we saw the Chernoff Bound, our first example of a concentration bound. Typically, this bound is used to show that a random variable is very close to its expected value with high probability. Such tools are extremely valuable to users of the probabilistic method as they allow us to show that with high probability, a random experiment behaves approximately as we “expect” it to.

The Chernoff Bound applies to a very limited type of random variable, essentially the number of heads in a sequence of tosses of the same weighted coin. While this limited tool can apply in a surprisingly large number of situations, it is often not enough. In this chapter, we will discuss Talagrand's Inequality, one of the most powerful concentration bounds commonly used in the probabilistic method. We will present another powerful bound, Azuma's Inequality, in the next chapter. These two bounds are very similar in nature, and both can be thought of as generalizations of the following:

Simple Concentration Bound *Let X be a random variable determined by n independent trials T_1, \dots, T_n , and satisfying*

changing the outcome of any one trial can affect X by at most c , (10.1)

then

$$\Pr(|X - \mathbf{E}(X)| > t) \leq 2e^{-\frac{t^2}{2c^2n}}.$$

Typically, we take c to be a small constant.

To motivate condition (10.1), we consider the following random variable which is *not* strongly concentrated around its expected value:

$$A = \begin{cases} n, & \text{with probability } \frac{1}{2} \\ 0, & \text{with probability } \frac{1}{2}. \end{cases}$$

To make A fit the type of random variable discussed in the Simple Concentration Bound, we can define T_1, \dots, T_n to be binomial random variables, each equal to 0 with probability $\frac{1}{2}$ and 1 with probability $\frac{1}{2}$, and set $A = 0$

if $T_n = 0$ and $A = n$ if $T_n = 1$. Here, $\mathbf{E}(A) = \frac{n}{2}$ but with probability 1, $|A - \mathbf{E}(A)| \geq \frac{n}{2}$. Contrast this with the random variable $B = \sum_{i=1}^n T_i$ (i.e. the number of 1's amongst T_1, \dots, T_n). The expected value of B is also $\frac{n}{2}$, but by the Chernoff Bound, the probability that $|B - \mathbf{E}(B)| \geq \alpha n$ is at most $2e^{-\frac{4\alpha^2 n}{3}}$. The difference is that B satisfies condition (10.1) with $c = 1$, while A clearly does not satisfy this condition unless we take c to be far too large to be useful. In essence, the outcomes of each of T_1, \dots, T_n combine equally to determine B , while A is determined by a single "all-or-nothing" trial. In the language of the stock markets, a diversified portfolio is less risky than a single investment.

It is straightforward to verify that the Simple Concentration Bound implies that $\Pr(|BIN(n, \frac{1}{2}) - \frac{n}{2}| > t) \leq 2e^{-\frac{t^2}{2n}}$, which is nearly as tight as the Chernoff Bound. (This is very good, considering that the Simple Concentration Bound is so much more widely applicable than the Chernoff Bound!) In the same way, for any constant p , the Simple Concentration Bound yields as good a bound as one can hope for on $\Pr(|BIN(n, p) - np| > t)$, up to the constant term in the exponent. However, when $p = o(1)$, the Simple Concentration Bound performs rather poorly on $BIN(n, p)$. For example, if $p = n^{-\frac{1}{2}}$, then it yields $\Pr(|BIN(n, p) - np| > \frac{1}{2}np) \leq 2e^{-\frac{1}{16}}$ which is far worse than the bound of $2e^{-\frac{\sqrt{n}}{12}}$ provided by the Chernoff Bound.

In general, we would often like to show that for any constant $\alpha > 0$ there exists a $\beta > 0$ such that $\Pr(|X - \mathbf{E}(X)| > \alpha \mathbf{E}(X)) \leq e^{-\beta \mathbf{E}(X)}$. The Simple Concentration Bound can only do this if $\mathbf{E}(X)$ is at least a constant fraction of n . Fortunately, when this property does not hold, Talagrand's Inequality will often do the trick.

Talagrand's Inequality adds a single condition to the Simple Concentration Bound to yield strong concentration even in the case that $\mathbf{E}(X) = o(n)$. In its simplest form, it yields concentration around the *median* of X , $\mathbf{Med}(X)$ rather than $\mathbf{E}(X)$. Fortunately, as we will see, if X is strongly concentrated around its median, then its expected value must be very close to its median, and so it is also strongly concentrated around its expected value, a fact that is usually much more useful as medians can be difficult to compute.

Talagrand's Inequality I *Let X be a non-negative random variable, not identically 0, which is determined by n independent trials T_1, \dots, T_n , and satisfying the following for some $c, r > 0$:*

1. *changing the outcome of any one trial can affect X by at most c , and*
2. *for any s , if $X \geq s$ then there is a set of at most rs trials whose outcomes certify that $X \geq s$ (we make this precise below),*

then for any $0 \leq t \leq \mathbf{Med}(X)$,

$$\Pr(|X - \mathbf{Med}(X)| > t) \leq 4e^{-\frac{t^2}{8c^2 r \mathbf{Med}(X)}}.$$

As with the Simple Concentration Bound, in a typical application c and r are small constants.

More precisely, condition 2 says that there is a set of trials T_{i_1}, \dots, T_{i_t} for some $t \leq rs$ such that changing the outcomes of all the other trials cannot cause X to be less than s , and so in order to “prove” to someone that $X \geq s$ it is enough to show her just the outcomes of T_{i_1}, \dots, T_{i_t} . For example, if each T_i is a binomial variable equal to 1 with probability p and 0 with probability $1 - p$, then if $X \geq s$ we could take T_{i_1}, \dots, T_{i_t} to be s of the trials which came up “1”.

The fact that Talagrand's Inequality proves concentration around the median rather than the expected value is not a serious problem, as in the situation where Talagrand's Inequality applies, those two values are very close together, and so concentration around one implies concentration around the other:

Fact 10.1 *Under the conditions of Talagrand's Inequality I, $|\mathbf{E}(X) - \text{Med}(X)| \leq 40c\sqrt{r\mathbf{E}(X)}$.*

This fact, which we prove in Chap. 20, now allows us to reformulate Talagrand's Inequality in terms of $\mathbf{E}(X)$.

Talagrand's Inequality II *Let X be a non-negative random variable, not identically 0, which is determined by n independent trials T_1, \dots, T_n , and satisfying the following for some $c, r > 0$:*

1. *changing the outcome of any one trial can affect X by at most c , and*
2. *for any s , if $X \geq s$ then there is a set of at most rs trials whose outcomes certify that $X \geq s$,*

then for any $0 \leq t \leq \mathbf{E}(X)$,

$$\Pr(|X - \mathbf{E}(X)| > t + 60c\sqrt{r\mathbf{E}(X)}) \leq 4e^{-\frac{t^2}{8c^2r\mathbf{E}(X)}}.$$

Remarks

1. The reason that the “40” from Fact 10.1 becomes a “60” here is that there is some loss in replacing $\text{Med}(X)$ with $\mathbf{E}(X)$ in the RHS of the inequality.
2. In almost every application, c and r are small constants and we take t to be asymptotically much larger than $\sqrt{\mathbf{E}(X)}$ and so the $60c\sqrt{r\mathbf{E}(X)}$ term is negligible. In particular, if the asymptotic order of t is greater than $\sqrt{\mathbf{E}(X)}$, then for any $\beta < \frac{1}{8c^2r}$ and $\mathbf{E}(X)$ sufficiently high, we have:

$$\Pr(|X - \mathbf{E}(X)| > t) \leq 2e^{-\frac{\beta t^2}{\mathbf{E}(X)}}.$$

3. This formulation is probably the simplest useful version of Talagrand's Inequality, but does not express its full power. In fact, this version *does not* imply the Simple Concentration Bound (as the interested reader may verify). In Chap. 20, we will present other more powerful versions of Talagrand's Inequality, including some from which the Simple Concentration Bound is easy to obtain. The version presented here is along the lines of a version developed by S. Janson, E. Shamir, M. Steele and J. Spencer (see [145]) shortly after the appearance of Talagrand's paper [148].

The reader should now verify that Talagrand's Inequality yields a bound on the concentration of $BIN(n, p)$ nearly as good as that obtained from the Chernoff Bound for all values of p .

To illustrate that Talagrand's Inequality is more powerful than the Chernoff Bound and the Simple Concentration Bound, we will consider a situation in which the latter two do not apply.

Consider a graph G . We will choose a random subgraph $H \subseteq G$ by placing each edge in $E(H)$ with probability p , where the choices for the edges are all independent. We define X to be the number of vertices which are endpoints of at least one edge in H .

Here, each random trial clearly affects X by at most 2, and so X satisfies condition (10.1). The problem is that if v is the number of vertices in G , then clearly $\mathbf{E}(X) \leq v$, while it is entirely possible that the number of edges in G , and hence the number of random trials, is of order v^2 . Thus, the Simple Concentration Bound does not give a good bound here (and the Chernoff Bound clearly doesn't apply). However, if $X \geq s$ then it is easy to find s trials which certify that X is at least s , namely a set of s edges which appear in H and which between them touch at least s vertices. Thus Talagrand's Inequality suffices to show that X is strongly concentrated.

We will present one final illustration, perhaps the most important of the simple applications of Talagrand's Inequality.

Let $\sigma = x_1, \dots, x_n$ be a uniformly random permutation of $1, \dots, n$, and let X be the length of the longest increasing subsequence of σ ¹. A well-known theorem of Erdős and Szekeres [47] states that any permutation of $1, \dots, n$ contains either a monotone increasing subsequence of length $\lceil \sqrt{n} \rceil$ or a monotone decreasing subsequence of length $\lceil \sqrt{n} \rceil$. It turns out that the expected value of X is approximately $2\sqrt{n}$, i.e. twice the length of a monotone subsequence guaranteed by the Erdős-Szekeres Theorem (see [104, 153]). A natural question is whether X is highly concentrated. Prior to the onset of Talagrand's Inequality, the best result in this direction was due to Frieze [61] who showed that with high probability, X is within a distance of roughly $\mathbf{E}(X)^{2/3}$ of its mean, somewhat weaker than our usual target of $\mathbf{E}(X)^{1/2}$.

At first glance, it is not clear whether Talagrand's Inequality applies here, since we are not dealing with a sequence of independent random trials.

¹ In other words, a subsequence $x_{i_1} < x_{i_2} < \dots < x_{i_k}$ where, of course, $i_1 < \dots < i_k$.

Thus, we need to choose our random permutation in a non-straightforward manner. We choose n uniformly random real numbers, y_1, \dots, y_n , from the interval $[0, 1]$. Now arranging y_1, \dots, y_n in increasing order induces a permutation σ of $1, \dots, n$ in the obvious manner².

If $X \geq s$, i.e. if there is an increasing subsequence of length s , then the s corresponding random reals clearly certify the existence of that increasing subsequence, and so certify that $X \geq s$. It follows that changing the value of any one y_i can affect X by at most one. So, Talagrand's Inequality implies:

$$\Pr(|X - \mathbf{E}(X)| > t + 60\sqrt{\mathbf{E}(X)}) < 4e^{-\frac{t^2}{8\mathbf{E}(X)}}.$$

This was one of the original applications of Talagrand's Inequality in [148]. More recently, Baik, Deift and Johansson [15] have shown that a similar result holds when we replace $\sqrt{\mathbf{E}(X)}$ by $\mathbf{E}(X)^{\frac{1}{3}}$, using different techniques (see also [96]).

We will find Talagrand's Inequality very useful when analyzing the Naive Colouring Procedure discussed in the introduction to this part of the book. To illustrate a typical situation, suppose that we apply the procedure using $\beta\Delta$ colours for some fixed $\beta > 0$, and consider what happens to the neighbourhood of a particular vertex v . Let A denote the number of colours assigned to the vertices in N_v , and let R denote the number of colours retained by the vertices in N_v after we uncolour vertices involved in conflicts.

If $\deg(v) = \Delta$, then it turns out that $\mathbf{E}(A)$ and $\mathbf{E}(R)$ are both of the same asymptotic order as Δ . A is determined solely by the colours assigned to the Δ vertices in N_v , and so a straightforward application of the Simple Concentration Bound proves that A is highly concentrated. R , on the other hand, is determined by the colours assigned to the up to Δ^2 vertices of distance at most two from v , and so the Simple Concentration Bound is insufficient here. However, as we will see, by applying Talagrand's Inequality we can show that R is also highly concentrated.

10.2 Colouring Triangle-Free Graphs

Now that we have Talagrand's Inequality in hand, we are ready to carry out our analysis of the naive random procedure we presented in the introduction to this part of the book. In this section we consider the special case of triangle-free graphs. We shall show:

Theorem 10.2 *There is a Δ_0 such that if G is a triangle-free graph with $\Delta(G) \geq \Delta_0$ then $\chi(G) \leq (1 - \frac{1}{2e^6})\Delta$.*

² Because these are uniformly random real numbers, it turns out that with probability 1, they are all distinct.

We shall improve significantly on this theorem in Chap. 13 where with the same hypotheses we obtain: $\chi(G) \leq O(\frac{\Delta}{\log \Delta})$.

Proof of Theorem 10.2. We will not specify Δ_0 , rather we simply insist that it is large enough so that certain inequalities implicit below hold. We can assume G is Δ -regular because our procedure from Sect. 1.5 for embedding graphs of maximum degree Δ in Δ -regular graphs maintains the property that G is triangle free.

We prove the theorem by finding a partial colouring of G using fewer than $\Delta - \frac{\Delta}{2e^6}$ colours such that in every neighbourhood there are at least $\frac{\Delta}{2e^6} + 1$ colours which appear more than once. As discussed in the introduction to this part of the book, we can then complete the desired colouring using a greedy colouring procedure (see Exercise 1.1). We find the required partial colouring by analyzing our naive random colouring procedure.

So, we set $C = \lfloor \frac{\Delta}{2} \rfloor$ and consider running our random procedure using C colours. That is, for each vertex w we *assign* to w a uniformly random colour from $\{1, \dots, C\}$. If w is assigned the same colour as a neighbour we uncolour it, otherwise we say w *retains* its colour.

We are interested, for each vertex v of G , in the number of colours which are assigned to at least two neighbours of v and retained on at least two of these vertices. In order to simplify our analysis, we consider the random variable X_v which counts the number of colours which are assigned to at least two neighbours of v and are retained by *all* of these vertices.

For each vertex v , we let A_v be the event that X_v is less than $\frac{\Delta}{2e^6} + 1$. We let $\mathcal{E} = \{A_v | v \in V(G)\}$. To prove the desired partial colouring exists, we need only show that with positive probability, none of these bad events occurs. We will apply the Local Lemma.

To begin, note that A_v depends only on the colour of vertices which are joined to v by a path of length at most 2. Thus, setting

$$S_v = \{A_w | v \text{ and } w \text{ are joined by a path of length at most } 4\},$$

we see that A_v is mutually independent of $\mathcal{E} - S_v$. But $|S_v| < \Delta^4$. So, as long as no A_v has probability greater than $\frac{1}{4\Delta^4}$, we are done.

We compute a bound on $\Pr(A_v)$ using a two step process which will be a standard technique throughout this book. First we will bound the expected value of X_v , and then we show that X_v is highly concentrated around its expected value:

Lemma 10.3 $\mathbf{E}(X_v) \geq \frac{\Delta}{e^6} - 1$.

Lemma 10.4 $\Pr(|X_v - \mathbf{E}(X_v)| > \log \Delta \times \sqrt{\mathbf{E}(X_v)}) < \frac{1}{4\Delta^5}$.

These two lemmas will complete our proof, because

$$\frac{\Delta}{e^6} - 1 - \log \Delta \times \sqrt{\mathbf{E}(X_v)} \geq \frac{\Delta}{e^6} - 1 - \log \Delta \sqrt{\Delta} > \frac{\Delta}{2e^6} + 1,$$

for Δ sufficiently large, and so these lemmas imply that $\mathbf{Pr}(A_v) < \frac{1}{4\Delta^5}$ as required.

Proof of Lemma 10.3. For each vertex v we define X'_v to be the number of colours which are assigned to exactly two vertices in $N(v)$ and are retained by both those vertices. Note that $X_v \geq X'_v$.

A pair of vertices $u, w \in N(v)$ will both retain the same colour α which is assigned to no other neighbour of v , iff α is assigned to both u and w and to no vertex in $S = N(v) \cup N(u) \cup N(w) - u - v$. Because $|S| \leq 3\Delta - 3 \leq 6C$, for any colour α , the probability that this occurs is at least $(\frac{1}{C})^2 \times (1 - \frac{1}{C})^{6C}$. There are C choices for α and $\binom{\Delta}{2}$ choices for $\{u, w\}$. Using Linearity of Expectation and the fact that $e^{-1/C} < 1 - \frac{1}{C} + \frac{1}{2C^2}$, we have:

$$\mathbf{E}(X'_v) \geq C \binom{\Delta}{2} \times \left(\frac{1}{C}\right)^2 \times \left(1 - \frac{1}{C}\right)^{6C} \geq \frac{\Delta - 1}{e^6} \left(1 - \frac{4}{C}\right) > \frac{\Delta}{e^6} - 1,$$

for C sufficiently large. \square

Proof of Lemma 10.4. Instead of proving the concentration of X_v directly, we will focus on two related variables. The first of these is AT_v (assigned twice) which counts the number of colours assigned to at least two neighbours of v . The second is Del_v (deleted) which counts the number of colours assigned to at least two neighbours of v but removed from at least one of them. We note that $X_v = AT_v - \text{Del}_v$, and so to prove Lemma 10.4, it will suffice to prove the following concentration bounds on these two related variables, which hold for any $t \geq \sqrt{\Delta \log \Delta}$.

$$\text{Claim 1: } \mathbf{Pr}(|AT_v - \mathbf{E}(AT_v)| > t) < 2e^{-\frac{t^2}{8\Delta}}.$$

$$\text{Claim 2: } \mathbf{Pr}(|\text{Del}_v - \mathbf{E}(\text{Del}_v)| > t) < 4e^{-\frac{t^2}{100\Delta}}.$$

To see that these two claims imply Lemma 10.4, we observe that by Linearity of Expectation, $\mathbf{E}(X_v) = \mathbf{E}(AT_v) - \mathbf{E}(\text{Del}_v)$. Therefore, if $|X_v - \mathbf{E}(X_v)| > \log \Delta \sqrt{\mathbf{E}(X_v)}$, then setting $t = \frac{1}{2} \log \Delta \sqrt{\mathbf{E}(X_v)}$, we must have either $|AT_v - \mathbf{E}(AT_v)| > t$ or $|\text{Del}_v - \mathbf{E}(\text{Del}_v)| > t$. Applying our claims, along with the Subadditivity of Probabilities, the probability of this happening is at most

$$2e^{-\frac{t^2}{8\Delta}} + 4e^{-\frac{t^2}{100\Delta}} < \frac{1}{4\Delta^5}.$$

It only remains to prove our claims.

Proof of Claim 1. The value of AT_v depends only on the Δ colour assignments made on the neighbours of v . Furthermore, changing any one of these assignments can affect AT_v by at most 2, as this change can only affect whether the old colour and whether the new colour are counted by AT_v . Therefore, the result follows from the Simple Concentration Bound with $c = 2$. \square

Proof of Claim 2. The value of Del_v depends on the up to nearly Δ^2 colour assignments made to the vertices of distance at most 2 from v . Because $\mathbf{E}(\text{Del}_v) \leq \Delta = o(\Delta^2)$ (since in fact Del_v is always at most Δ), the Simple Concentration Bound will not apply here. So we will use Talagrand's Inequality.

As with AT_v , changing any one colour assignment can affect Del_v by at most 2. Furthermore, if $\text{Del}_v \geq s$ then there is a set of at most $3s$ colour assignments which certify that $\text{Del}_v \geq s$. Namely, for each of s colours counted by Del_v , we take 2 vertices of that colour in N_v and one of their neighbours which also received that colour. Therefore, we can apply Talagrand's Inequality with $c = 2$ and $r = 3$ to obtain:

$$\Pr(|\text{Del}_v - \mathbf{E}(\text{Del}_v)| > t) < 4e^{-\frac{(t - 120\sqrt{3\mathbf{E}(\text{Del}_v)})^2}{96\mathbf{E}(\text{Del}_v)}} < 4e^{-\frac{t^2}{100\Delta}},$$

since $t \geq \sqrt{\Delta \log \Delta}$ and $\mathbf{E}(\text{Del}_v) \leq \Delta$. □

So, we have proven our two main lemmas and hence the theorem. □

10.3 Colouring Sparse Graphs

In this section, we generalize from graphs in which each neighbourhood contains no edges to graphs in which each neighbourhood contains a reasonable number of non-edges. In other words, we consider graphs G which have maximum degree Δ and such that for each vertex v there are at most $\binom{\Delta}{2} - B$ edges in the subgraph induced by $N(v)$ for some reasonably large B .

We note that to ensure that G has a $\Delta - 1$ colouring we will need to insist that B is at least $\Delta - 1$ as can be seen by considering the graph obtained from a clique of size Δ by adding a vertex adjacent to one element of the clique. More generally, we cannot expect to colour with fewer than $\Delta - \lfloor \frac{B}{\Delta} \rfloor$ colours for values of B which are smaller than $\Delta^{\frac{3}{2}} - \Delta$ (see Exercise 10.1). We shall show that if B is not too small, then we can get by with nearly this small a number of colours.

Theorem 10.5 *There exists a Δ_0 such that if G has maximal degree $\Delta > \Delta_0$ and $B \geq \Delta(\log \Delta)^3$, and no $N(v)$ contains more than $\binom{\Delta}{2} - B$ edges then $\chi(G) \leq \Delta + 1 - \frac{B}{e^6 \Delta}$.*

Proof The proof of this theorem mirrors that of Theorem 10.2. Once again, we will consider X_v , the number of colours assigned to at least two *non-adjacent* neighbours of v and retained on all the neighbours of v to which it is assigned. Again, the proof reduces to the following two lemmas, whose (omitted) proofs are nearly identical to those of the corresponding lemmas in the previous section.

Lemma 10.6 $\mathbf{E}(X_v) \geq \frac{2B}{e^6 \Delta}.$

Lemma 10.7 $\Pr\left(|X_v - \mathbf{E}(X_v)| > \log \Delta \sqrt{\mathbf{E}(X_v)}\right) < \frac{1}{4\Delta^5}.$

It is easy to see that provided B is at least $\Delta(\log \Delta)^3$ then we can combine these two facts to obtain the desired result. \square

Remarks

1. Using much more complicated techniques, such as those introduced in Part 6, we can show that Theorem 10.5 holds for every B .
2. As you will see in Exercise 10.2, a slight modification to this proof yields the same bound on the list chromatic number of G .

Using our usual argument, we can obtain a version of Theorem 10.5 which holds for every Δ , by weakening our constant e^6 :

Corollary 10.8 *There exists a constant $\delta > 0$ such that if G has maximal degree $\Delta > 0$ and $B \geq \Delta(\log \Delta)^3$, and no $N(v)$ contains more than $\binom{\Delta}{2} - B$ edges then $\chi(G) \leq \Delta + 1 - \delta \frac{B}{\Delta}.$*

Proof of Corollary 10.8. Set $\delta = \min\{\frac{1}{\Delta_0}, \frac{1}{e^6}\}$. Now, if $\Delta > \Delta_0$ then the result holds by Theorem 10.5. Otherwise it holds because $B < \frac{\Delta}{\delta}$ so we are simply claiming that G is Δ colourable, which is true by Brooks' Theorem. \square

10.4 Strong Edge Colourings

We close this chapter by describing an application of Theorem 10.5 to strong edge colourings, which appears in [118] and which motivated Theorem 10.5. The notion of a strong edge colouring was first introduced by Erdős and Nešetřil (see [52]), and is unrelated to the strong vertex colourings which were discussed in Chap. 8.

A *strong edge-colouring* of a graph, G , is a proper edge-colouring of G with the added restriction that no edge is adjacent to two edges of the same colour, i.e. a 2-frugal colouring of $L(G)$ (note that in any proper edge-colouring of G , no edge is adjacent to three edges of the same colour and so the corresponding colouring of $L(G)$ is 3-frugal). Equivalently, it is a proper vertex-colouring of $L(G)^2$, the square of the line graph of G which has the same vertex set as $L(G)$ and in which two vertices are adjacent precisely if they are joined by a path of length at most two in $L(G)$. The *strong chromatic index* of G , $s\chi_e(G)$ is the least integer k such that G has a strong edge-colouring using k colours. We note that each colour class in a strong edge colouring is an *induced matching*, that is a matching M such that no edge of $G - M$ joins endpoints of two distinct elements of M .




Fig. 10.1. An edge and its neighbourhood in $L(G)^2$

If G has maximum degree Δ , then trivially $s\chi_e(G) \leq 2\Delta^2 - 2\Delta + 1$, as $L(G)^2$ has maximum degree at most $2\Delta^2 - 2\Delta$. In 1985, Erdős and Nešetřil pointed out that the graph G_k obtained from a cycle of length five by duplicating each vertex k times (see Fig. 10.2) contains no induced matching of size 2. Therefore, in any strong colouring of G_k , every edge must get a different colour, and so

$$s\chi_e(G_k) = |E(G_k)| = 5k^2.$$

We note that $\Delta(G_k) = 2k$, and so $s\chi_e(G_k) = \frac{5}{4}\Delta(G_k)^2$.

Erdős and Nešetřil conjectured that these graphs are extremal in the following sense:

Conjecture For any graph G , $s\chi_e(G) \leq \frac{5}{4}\Delta(G)^2$.

They also asked if it could even be proved that for some fixed $\epsilon > 0$ every graph G satisfies $s\chi_e(G) \leq (2 - \epsilon)\Delta(G)^2$. In this section we briefly point out




Fig. 10.2. G_3

how Theorem 10.5 can be used to answer this question in the affirmative. For other work on this and related problems, see [11], [33], [53], [83] and [84].

Theorem 10.9 *There is a Δ_0 such that if G has maximum degree $\Delta \geq \Delta_0$, then $s\chi_e(G) \leq 1.99995\Delta^2$.*

Corollary 10.10 *There exists a constant $\epsilon > 0$ such that for every graph G , $s\chi_e(G) \leq (2 - \epsilon)\Delta^2$.*

Proof of Corollary 10.10. Set $\epsilon = \min\{.00005, \frac{1}{\Delta_0}\}$. If $\Delta(G) \geq \Delta_0$ then the result follows from the theorem. Otherwise, we simply need to use the fact that $\Delta(L(G)^2) \leq 2\Delta(G)^2 - 2\Delta(G)$ and apply Brook's Theorem.

To prove Theorem 10.9, the main step is to show the following, whose tedious but routine proof can be found in [118]:

Lemma 10.11 *If G has maximum degree Δ sufficiently large then, for each $e \in V(L(G)^2)$, $N_{L(G)^2}(e)$ has at most $(1 - \frac{1}{36})\binom{2\Delta^2}{2}$ edges.*

Using this lemma, it is a straightforward matter to apply Theorem 10.5 to yield Theorem 10.9.

Exercises

Exercise 10.1 Show that for every Δ and $B \leq \Delta^{\frac{3}{2}} - \Delta$ there exists a graph with $\Delta + 1$ vertices in which the neighbourhood of each vertex contains at most $\binom{\Delta}{2} - B$ edges, whose chromatic number is $\lfloor \Delta - \frac{B}{\Delta} \rfloor$.

Exercise 10.2 Complete the proof of Theorem 10.5 and then modify it to get the same bound on the list chromatic number of G .



<http://www.springer.com/978-3-540-42139-9>

Graph Colouring and the Probabilistic Method

Molloy, M.; Reed, B.

2002, XIV, 326 p., Hardcover

ISBN: 978-3-540-42139-9