Graph Colouring and the Rödl Nibble

Ziyang Jin[1]

[1]Theory Group
Department of Computer Science
University of Toronto

11 October 2023

Table of Contents

What Kind of Colouring

When you work on graph colouring, the usual question is **what kind of colouring** can be achieved on **what kind of graph** (using the probability method)?

What kind of colouring? Vertex colouring, edge colouring, list colouring, frugal colouring, ...

What kind of graph? Any graph, sparse graph, dense graph, triang-free graph, graph with girth at least five...

Lovász Local Lemma

Theorem (Lovász Local Lemma)

Consider a set $\mathcal{E}$ of (typically bad) events such that for each $A \in \mathcal{E}$:
(a) $\Pr(A) \leq p < 1$, *and*
(b) A is mutually independent of a set of all but at most d of the other events
If $4pd \leq 1$, *then with positive probability, none of the events in* $\mathcal{E}$ *occur.*

Note that $4pd$ can be replaced by epd and the theorem still holds.

Mutual Independence Principle

Theorem (The Mutual Independence Principle)

Suppose that $\mathcal{X} = X_1, \ldots, X_m$ is a sequence of independent random experiments. Suppose further that $A_1, \ldots, A_n$ is a set of events, where each A_i is determined by $F_i \subseteq \mathcal{X}$. If $F_i \cap (F_{i_1}, \ldots, F_{i_k}) = \emptyset$ then A_i is mutually independent of $\{A_{i_1}, \ldots, A_{i_k}\}$.

To put it in simple words, if some event A is determined by some random experiments $X_1, \ldots, X_m$, then A is independent from the other events that are not determined by any of $X_1, \ldots, X_m$.

Consider the Neighbourhood

When we want to prove some colouring property using the Loász Local Lemma, we consider a "local" part of the graph, i.e., an arbitrary vertex v and its neighbourhood $u_1, ..., u_k \in N(v)$.

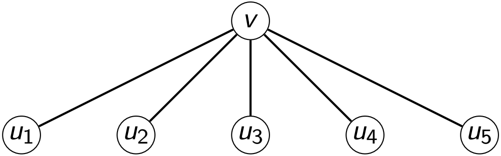

If any arbitrary neighbourhood is "good", then by Loász Local Lemma, the whole graph is "good". You can think of "good" as, for example, no edges have conflicting colours, etc.

The Problem Statement

We would like to prove the following theorem.

Theorem (Kim '95)

We can colour any graph with girth at least five using $(1+o(1))\frac{\Delta}{\ln \Delta}$ *colours, where* Δ *is the maximum degree of the graph.*

Definition

The **girth** of a graph is the length of a shortest cycle contained in the graph.

The Intuition

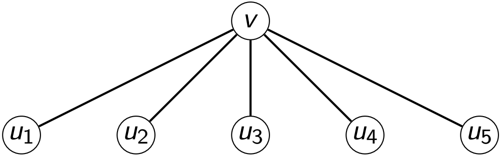

We could like to colour the graph using C colours. For any vertex v, for each neighbour $u \in N(v)$, let the list of available colours $L_u = \{1, 2, ..., C\}$. We independently assign each $u \in N(v)$ a colour drawn uniformly at random from L_u. Suppose $|N(v)| = \Delta$. What is the expected number of colours available to v?

Calculation

Let X_i be the random variable that colour i is available to v, i.e. colour i is not assigned to any of the vertices in the neighbourhood.

$$\Pr(X_i) = (1 - \frac{1}{C})^{\Delta}$$

Let $X = \sum_{i=1}^{C} X_i$. By linearity of expectation, we have

$$\mathbb{E}(X) = \sum_{i=1}^{C} \Pr(X_i) = C(1 - \frac{1}{C})^{\Delta} \approx Ce^{-\Delta/C}$$

Let $Ce^{-\Delta/C} \geq 1$, we can get $C \geq (1+\epsilon)\frac{\Delta}{\ln \Delta}$, and the number of colours available to v will be $(1+\epsilon)\frac{\Delta^{\epsilon/1+\epsilon}}{\ln \Delta}$, which goes to infinity as $\Delta \to \infty$.

Relate Intuition to Girth Five Graph

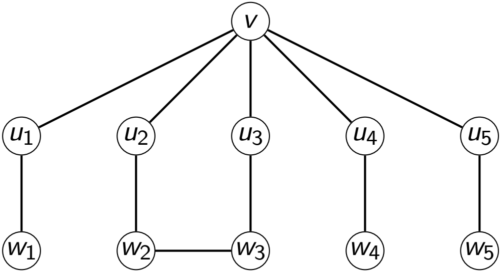

Consider any vertex v and its neighbourhood. For any two neighbours $u, u' \in N(v)$, the only common neighbour u and u' share is v. Suppose in the very beginning, every vertex has a list of colours $\{1, ..., C\}$. Suppose all the vertices outside of $v \cup N(v)$ have been coloured. What can we say about L_u and L'_u?

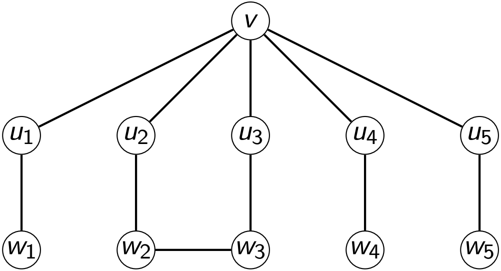

L_u ($L_{u'}$) is the list of colours available to u (u').

$L_u = \{1, .., C\} -$ set of colours assigned to $N(u)$.

$L_{u'} = \{1, .., C\} -$ set of colours assigned to $N(u')$.

Since u and u' do not share any neighbours other than v, L_u and L'_u should look **almost independent**.

Note that we cannot say the same for triangle-free graphs, i.e., graphs with girth at least four.

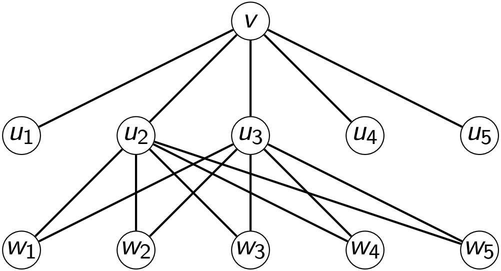

If u and u' have many common neighbours, then

$$\text{set of colours assigned to } N(u) \approx \text{set of colours assigned to } N(u')$$

which means L_u and $L_{u'}$ will be (strongly) dependent.

Exercise

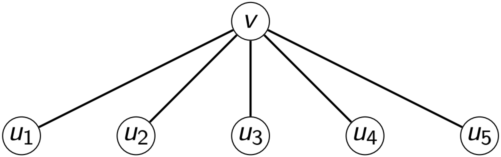

If we want to get an intuition for the chromatic number of triangle-free graphs, what kind of random analysis should we perform, i.e., what should we assume on the list of available colour L_u for each $u \in N(v)$?

Hint: For analyzing graphs with girth at least five, we independently assign each $u \in N(v)$ a colour randomly from L_u, and each $L_u = \{1, ..., C\}$, and compute the expected number of colours available to v.

The Rödl Nibble and Colouring Graphs with Girth At Least Five

The Dice Analogy

Suppose I throw 1000 dices, and I would like to know if I can get an instance where the mean is exactly $(1+6)/2 = 3.5$.

Assume the dices are independent, we know this is certainly possible, but how many rounds do I need play in order to get the mean to be exactly 3.5 (not $3.5 \pm \epsilon$)?

A lot of rounds!

What if I throw two dices at a time, and once their mean is 3.5, fix them, and keep throwing another two, then fix, then another two, ..., until I finish all 1000 dices.

Obviously, it is easier to achieve an instance where the mean is exactly 3.5 using this approach.

This is the idea of **the Rödl Nibble**. If we want to get an instance with some specific property, instead of performing random experiment on the whole set of objects, we do it bite by bite (the "nibble"), i.e., we perform random experiments on a subset of objects, and once we get the property we want, fix this subset, and move on to another set of objects.

The Bite Size

Why do not we perform random experiment on one object at a time?

Usually random experiment on one object do not give you concentration. So we want to carefully choose the "bite size" such that we have both concentration and it is easy to achieve.

The Rödl nibble is also called the semi-random method. It is named after Vojtěch Rödl (Czech), who used this approach to solve packing in a hypergraph problem in response to the Erdös-Hanani conjecture.

Currently, there is no official Wikipedia definition for the Rödl nibble or the semi-random method. It is currently a subsection under the Wikipedia page “Packing in a hypergraph”, but its has wide applications in combinatorics.

We are not going to show the original proof for hypergraph packing. Instead, we will see an application in colouring graphs with girth at least five.

The Colouring Procedure

Wasteful Colouring Procedure

1. For each uncoloured vertex v, activate it with probability $\frac{K}{\ln \Delta}$.
2. For each activated vertex v, assign to it a random colour from L_v.
3. For each activated vertex v, removed the colour assigned to v from L_u for $u \in N(v)$.
4. Uncolour every vertex which receives the same colour as a neighbour simultaneously.
5. Conduct an "equalizing" coin flip for each vertex v and colour $c \in L_v$, and remove c if it loses the coin flip.

- K is a small constant to be decided later.
- The last step is unnecessary, but it makes the analysis much easier.

UNIVERSITY OF TORONTO

The **wastefulness** comes from:

- Once we assigned the colour c to v, we remove c from L_u for all $u \in N(v)$. However, if c is uncoloured. We do not give c back to the L_u's.
- The equalizing coin flip.

Luckily, we are not wasting too much, so we can still complete the colouring even though we throw out some colours during the process.

One reason is that $(1+\epsilon)\frac{\Delta}{\ln \Delta}$ is not the absolute low-bound. The absolute lower bound is $(\frac{1}{2}+\epsilon)\frac{\Delta}{\ln \Delta}$, but in order to get anything below 1, you need an analysis beyond looking at a local neighbourhood.

Open Question: How to analyse beyond looking at a local neighbourhood? What tools can be used?

The **Wasteful Colouring Procedure** is one iteration of our colouring algorithm.

Our algorithm starts with an uncoloured graph, and $L_v = \{1, ..., C\}$ for every $v \in V$, then it runs the Wasteful Colouring Procedure, and checks the (partial) colouring after the iteration. If some certain properties are satisfied, then we continue on executing the next iteration. Otherwise, we undo the current iteration and repeat the current iteration until the properties are satisfied.

We stop running the iterations when some termination criteria is met, and we finish the colouring under the termination condition. This is usually called the "finishing blow" in graph colouring.

The Finishing Blow

We stop running the iterations when for every vertex v, L_v has at least ℓ colours, and for every colour $c \in L_v$, it is acceptable for at most $\frac{\ell}{8}$ neighbours of v.

Claim. We can complete the colouring under the above condition.

This can be proved by a simple Loász Local Lemma argument. For each remaining uncoloured vertex v, we independently assign to v a random colour from its list of available colours L_v.

UNIVERSITY OF TORONTO

Proof.

Let $A_{e,i}$ be the bad event that edge $e = (u, v)$ gets assigned colour i on both endpoints. So $\Pr(A_{e,i}) = \frac{1}{\ell^2}$. What is the number of dependent events d to $A_{e,i}$? By principal of mutual independence, $A_{e,i}$ is determined by the colour assignment on u and v. Therefore, d is at most $2 \times \ell \times \frac{\ell}{8} = \frac{\ell^2}{4}$. Therefore, the LLL condition $4pd \leq 4 \times \frac{1}{\ell^2} \times \frac{\ell^2}{4} \leq 1$ is satisfied, there is positive probability that none of $A_{e,i}$ happens, which means the graph has a proper colouring. □

How do we transform the finishing blow to an algorithm?

Answer: **Entropy compression!** [Moser '08][Moser,Tardos '10] (See my last TSS talk.)

Basically, you can convert almost any L.L.L. argument proof into a randomized algorithm that terminates with high probability.

Property

What property do we need to satisfy in order to enter the next iteration?

The terminating condition – for every vertex v, L_v has at least ℓ colours, and for every colour $c \in L_v$, it is acceptable for at most $\frac{\ell}{8}$ neighbours of v.

Define $l_i(v)$ to be the number of available colour to v at the beginning of iteration i.
Define $t_i(v, c)$ to be the number of uncoloured neighbours u of v such that $c \in L_u$ at the beginning of iteration i.

Define **Property P(i)**:

$$l_i(v) \geq L_i$$

$$t_i(v, c) \leq T_i$$

where L_i and T_i are recursively defined throughout the iterations.

If some colour c is assigned to vertex v during the Wasteful Colouring Procedure, let $Keep_i(v, c)$ be the probability that c is maintained in vertex v at the end of the procedure, i.e., v is not uncoloured due to a conflict, and c is also not removed due to losing an equalizing coin flip.

$$Keep_i(v, c) = \prod_{u \in N(v), c \in L_u} (1 - \frac{K}{\ln \Delta} \times \frac{1}{l_i(u)})$$

which is greater than or equal to

$$Keep_i = (1 - \frac{K}{\ln \Delta} \times \frac{1}{L_i})^{T_i}$$

Essentially, the equalizing coin flip is to make every $Keep_i(v, c) = Keep_i$ to simplify the analysis. So we will just use $Keep_i$ from now on.

Expectations of l_{i+1} and t_{i+1}

Let's compute the expectations of $l_{i+1}(v)$ and $t_{i+1}(v, c)$.

$$\mathbb{E}[l_{i+1}(v)] = l_i(v) \times Keep_i$$

$$\mathbb{E}[t_{i+1}(v, c)] \approx t_i(v, c)(1 - \frac{K}{\ln \Delta} \times Keep_i) \times Keep_i$$

The first $Keep_i$ inside the parenthesis is essentially $Keep_i(u, c)$ for $u \in N(v)$, i.e., the probability that an activated neighbour retains its colour. The second $Keep_i$ outside is essentially $Keep_i(v, c)$, i.e., the probability that colour c itself is retained in L_v at the end of iteration i. We use the approximation sign as the two events are not independent, but close enough.

The recursive definition of L_i and T_i are almost the same as the expectation $\mathbb{E}[l_{i+1}(v)]$ and $\mathbb{E}[t_{i+1}(v,c)]$. Set $L_1 = (1+\epsilon)\frac{\Delta}{\ln \Delta}$ and $T_1 = \Delta$, and we recursively define

$$L_{i+1} = L_i \times Keep_i - L^{2/3}$$

$$T_{i+1} = T_i(1 - \frac{K}{\ln \Delta} \times Keep_i) \times Keep_i + T^{2/3}$$

The $L^{2/3}$ and $T^{2/3}$ here is that if we set L_i and T_i to be the exact expectation, it is difficult to prove an instance where $l_i(v) \geq \mathbb{E}[l_{i+1}(v)]$ and $t_i(v,c) \leq \mathbb{E}[t_{i+1}(v,c)]$ is true for every vertex in the graph simultaneously. So we make some relaxation here.

What to prove

1. The expectation calculation is correct.

Lemma (1: expectation)

$$\mathbb{E}[l_{i+1}(v)] = l_i(v) \times Keep_i$$

$$\mathbb{E}[t_{i+1}(v,c)] = t_i(v,c)(1 - \frac{K}{\ln \Delta} \times Keep_i) \times Keep_i + O(\frac{T_i}{L_i})$$

UNIVERSITY OF TORONTO

2. The random variables are concentrated around the expectation.

Lemma (2: concentration)

$$\Pr(|l_{i+1}(v) - \mathbb{E}(l_{i+1}(v))| > L^{2/3}) < \Delta^{-\ln \Delta}$$

$$\Pr(|t'_{i+1}(v,c) - \mathbb{E}(t'_{i+1}(v,c))| > \frac{1}{2}T^{2/3}) < \Delta^{-\ln \Delta}$$

We use $t'_{i+1}(v,c)$ here because $t_{i+1}(v,c)$ is not concentrated as if v is assigned colour c during the iteration, it removes c from all neighbours colouring list, i.e., $t_{i+1}(v,c) = 0$. $t'_{i+1}(v,c)$ is carefully worded to avoid such situation.

Note that we only use "the graph has girth at least five" condition in proving the concentration bounds, where we use Talagrand's inequality. We will save the discussion on Talagrand's inequality for next talk!

3. Using concentration and a simple L.L.L argument, we can show that **Property P(i)** holds for every iteration before the termination condition.

Lemma (3: P(i) holds for every iteration)

With positive probability, P(i) holds for every i such that for all $1 \leq j < i : L_j, T_j \geq \ln^7 \Delta$ *and* $T_j \geq \frac{1}{8} L_j$.

In the previous lemma, we have proved the probability $p \leq \Delta^{-\ln \Delta}$, and the number of dependent events (using Principle of Mutual Independence) $d = O(\Delta^5)$. So the L.L.L. condition $4pd \leq 4\Delta^{-\ln \Delta}\Delta^5 \leq 1$ is satisfied for large enough Δ.

4. Although relaxed, our definition of L_i and T_i is not so different from $\mathbb{E}[l_{i+1}(v)]$ and $\mathbb{E}[t_{i+1}(v, c)]$.

Lemma (4: L_i, T_i close to expectation)

Define

$$L'_{i+1} = L'_i \times Keep_i$$

$$T'_{i+1} = T'_i(1 - \frac{K}{\ln \Delta} \times Keep_i) \times Keep_i$$

If for all $1 \leq j < i$ *we have* $L_j, T_j \geq \ln^7 \Delta$ *and* $T_j \geq \frac{1}{8} L_j$, *then*
(a) $|L_i - L'_i| \leq (L'_i)^{5/6}$;
(b) $|T_i - T'_i| \leq (T'_i)^{5/6}$.

Note that $(L'_i)^{5/6} = o(L'_i)$ and $(T'_i)^{5/6} = o(T'_i)$.

5. In the end, there exists an iteration i^* such that we meet the termination condition at iteration i^*.

Lemma (5: can achieve termination condition)

There exists i^* such that:
(a) For all $i \leq i^*$, $T_i \geq \ln^8 \Delta$, $L_i > \Delta^{\epsilon/3}$, and $T_i \geq \frac{1}{8} L_i$;
(b) $T_{i^*+1} \leq \frac{1}{8} L_{i^*+1}$.

Proof of Main Theorem

Proof.

In the beginning, the graph is uncoloured and every vertex v has a list of $|L_v| = (1+\epsilon)\frac{\Delta}{\ln \Delta}$ colours. We carry out our Wasteful Colouring Procedure for up to i^* iterations. If **P(i)** fails to hold for any iteration i, then we halt. By Lemma 3 (**P(i)** holds for every iteration) and Lemma 5 (can achieve termination condition), with positive probability **P(i)** holds for every iteration and so we do in fact perform i^* iterations. After iteration i^*, we meet the termination condition so we can apply the finishing blow to complete the colouring. □

Q & A

Questions?

Thank you

Thank you!