

Class 12: Agenda and Questions

1 Announcements

- HW6 due Wednesday!
- You get a week off before HW7! (Due to President's day)

2 Questions?

Any questions from the minilectures and/or the quiz? (Constructive LLL)

3 You prove the constructive LLL for another problem!

Consider the following problem (which has featured on a quiz). You are coloring the integers $\{1, \dots, n\}$ either blue or red. You are given as input a collection of sets $S_1, S_2, \dots, S_m \subseteq \{1, \dots, n\}$, so that:

- Each set S_i has size at most k .
- Each set S_i intersects at most d other sets S_j , for some $d > 1$.

Our goal is to color the points $\{1, \dots, n\}$ so that there is no monochromatic set S_i . (A set S_i is *monochromatic* if every element of it is either red or blue).

Group Work

1. Mimic the proof of the constructive LLL that we saw for k -SAT to give a randomized algorithm that does the following.

Suppose that $k \geq \log_2 d + 10000$. (Here, 10000 is a stand-in for “some big enough constant.”) Then there is a randomized algorithm that proceeds by re-randomizing the sets S_i , (that is, it will iteratively look at different sets S_i and randomly re-color all of the points in that set), so that:

- If the algorithm terminates, then all of the numbers $\{1, \dots, n\}$ will be colored so that there is no monochromatic set S_j .
- The expected number of times that the algorithm re-randomizes a set S_j is $\text{poly}(m)$.

Don't worry about giving a complete proof with all the details, just work it out with enough detail that you believe it. As we did in the minilecture video, you may

use the (informal) fact that “there is no function $f : \{0, 1\}^X \rightarrow \{0, 1\}^Y$ so that (a) $Y \ll X$ and (b) with high probability over a uniformly random $x \in \{0, 1\}^X$, it is possible to recover x given $f(x)$.”

Hint: To map this problem onto k -SAT, think of the S_j 's as standing in for clauses, and the numbers $\{1, \dots, n\}$ as standing in for variables.

Hint: It's not quite as straightforward as applying the mapping in the previous hint and calling it a day. In particular, can you still work backwards from the “print” statements in the k -SAT version to figure out the original random bits?

2. What happens to your proof if the number of possible colors grows from two (blue and red) to some number t ? In particular, can you get the same guarantee as above, but under a weaker guarantee (eg, $k \geq \lceil \text{something smaller than } \log d + 10000 \rceil$).
3. How does the answer that you got in the previous part compare to what Corollary 3 in the lecture notes would give you for this problem?

As a reminder, that Corollary says:

Let V be a finite set of independent random variables. Let \mathcal{A} be a finite set of events determined by the random variables in V . If for all $A \in \mathcal{A}$, $|\Gamma(A)| \leq d + 1$, and $\Pr[A] \leq \frac{1}{e(d+1)}$, then the algorithm from the lecture notes (the general one, not just for k -SAT) will find an assignment to the variables V such that no event of \mathcal{A} occurs. Additionally, the expected number of “re-randomizations” performed by the algorithm is bounded by $O(|\mathcal{A}|/(d + 1))$.

4. **[This question is open-ended and may be difficult—think about it after you finish the others if you still have time.]** What happens to your proof if the sets can have variable size? (e.g., if all but a few of them have size k , and a few can be really small? Or if they have average size k ? Or....?)

Group Work: Solutions

1. Here is our algorithm with print statements that will help in the proof. We've just copy-pasted the algorithm from the lecture notes (made slightly informal for ease of discussion), but changed all the talk of satisfying formulas to coloring sets. The only change is the line in red: we needed to give more information about what went wrong.

Algorithm 1. FIND ASSIGNMENT AND PRINT STUFFGiven S_1, \dots, S_m :

- Choose a random assignment σ that assigns a color to each of the numbers $1, \dots, n$.
- For each set S_i that is monochromatic under σ :
 - print "Running FIX on Set i "
 - Run $\text{FIX}(i, \sigma)$ on S_i and σ .
- Return σ .

Algorithm 2. $\text{FIX}(i, \sigma)$ Given S_1, \dots, S_m , i , and σ :

- If we have run for too long (T re-randomizations), print "I give up; I've got σ " and halt.
- Re-randomize the colors of all of the variables in S_i .
- Suppose that $S_{i_1}, \dots, S_{i_{d+1}}$ are the sets that intersect with S_i (including S_i itself).
- For $\ell = 1, 2, \dots, d+1$:
 - If S_{i_ℓ} is monochromatic (say all [color], where [color] is either blue or red):
 - * Print "Trying to fix the ℓ 'th child..."
 - * Print "the problem was that it was all [color]"
 - * Run $\text{FIX}(i_\ell, \sigma)$
- print "All done, moving back up a level."

To do the proof, we need to do the following three steps:

- A We need to argue that we can recover the random bits that go in from the print statements that come out.
- B We need to count the number of random bits that go in to our algorithm in T re-randomizations (say in one call to FIX)
- C We need to count the number of bits that come out in T re-randomizations (in one call to FIX).

Thing A is why we needed to add the extra info about the color. With that, let's see why A is true. Suppose that we are working backwards from the final σ returned. Inductively assume that we know the assignment *after* the t 'th re-randomization, and we want to know the assignment σ' *before* that randomization, and we want to know the outcome of the random bits used in the t 'th re-randomization. We know which set S_i was re-randomized at the t 'th step, since we are using the "fixing ℓ 'th child" statements to follow along through the execution tree. Thus, every number that was *not* in S_i didn't change, so we know what colors all of those were in σ' . For the items in S_i , our extra print statement told us that before they were all monochromatic [blue or red], so we know what those should have been in σ' as well,

and we know what the randomizing bits were. Continuing in this way, we can work backwards all the way to the original assignment σ , and we can recover the random bits that were flipped.

For B, the number of random bits we use is $n + Tk$, since there are n bits for the original assignment, and k for each of the T re-randomizations.

For C, the number of bits of print statements that come out are:

- $m(C + \log m)$ at the beginning of highest-level calls to `FIX` to say which clause we are starting on.
- $T(\log(d+1)+1+C)$ bits for “trying to fix the ℓ 'th child because it was red” sorts of calls. (Notice that “because it was red” is only one extra bit of information).
- $n + C$ bits for “I give up, here's σ ”

where above C is some constant. Thus, the total is at most

$$m(\log m + C) + T(\log(d + 1) + C) + n + C$$

Now, we put it all together. In order to get a contradiction (because the number of bits out would be way less than the number in), we want

$$n + kT \gg m(\log m + C) + T(\log(d + 1) + C) + n + C.$$

After some algebra, we want

$$m(\log m + C) \ll T(k - \log(d + 1) + C)$$

Thus, provided that $k \geq \log(d+1) + C$ for some constant C , the RHS is a constant, and we win when $T \approx m \log m$. This is what we wanted to show.

2. If the number of colors grows from 2 to t , then we just have to change the algorithm to say “because it was green” or “because it was purple” or whatever. This is $\log t$ bits of communication. Working that into our computations above, we get that we need

$$n \log t + kT \log t \gg m(\log m + C) + T(\log(d + 1) + \log(t) + C) + n \log t + C.$$

aka

$$m(\log m + C) \ll T(k \log t - \log(d + 1) + \log t + C)$$

Thus, we'd need $k \geq \log_t(d + 1) + C$ for some constant C , and again we get $O(m \log m)$ re-randomizations.

3. In the context of Cor. 3, we have $V = \{1, \dots, n\}$ and $\mathcal{A} = \{A_i\}_{i \in [m]}$ where A_i is the event that S_i is monochromatic. If there are t colors, then the probability that A_i is monochromatic is $t^{-(k-1)}$. Thus, Cor. 3 says we need

$$t^{-(k-1)} \leq 1/(e(d+1)),$$

or that $k \geq \log_t(e(d+1)) + 1$, or that

$$k \leq \log_t(d+1) + C,$$

where $C = 1 + \log_t(e)$. This is roughly the same as what we got above (maybe with a different additive constant).

4 Using the asymmetric LLL

(We will (probably) not cover this in class; it's here in case there's extra time and/or you finish early)

Recall that the asymmetric LLL (that is, the more general statement that we had in the lecture notes, both for the algorithmic and non-algorithmic version) was:

Theorem 1. *Let V be a finite set of independent random variables. Let \mathcal{A} be a finite set of events determined by the random variables in V . If there exists an assignment $x : \mathcal{A} \rightarrow (0, 1)$ such that for all $A \in \mathcal{A}$,*

$$\Pr[A] \leq x(A) \prod_{B \in \Gamma(A) \setminus \{A\}} (1 - x(B)),$$

then [the algorithm from the mini-lecture] will find an assignment to the variables V such that no event of \mathcal{A} occurs. Additionally, the expected number of "re-randomizations" is bounded by $\sum_{A \in \mathcal{A}} \frac{x(A)}{1-x(A)}$.

We saw in the mini-lecture how to derive the symmetric version from this (set $x(A) = 1/(d+1)$ for all A), but it might seem pretty unclear how to apply it in the general asymmetric case. In this exercise you'll work out an example.

Group Work

Suppose that n and ℓ are sufficiently large, and suppose that $\ell \gg \sqrt{n} \cdot \log(n)$. Show that there is an edge-coloring of the complete graph on n vertices by two colors (blue or red) that contains *neither* a blue triangle *nor* a red ℓ -clique. (And that the algorithm from class will find such a coloring).

(Don't worry too much about constants in your proof; in particular, it's fine to treat $e^x \approx 1 + x$ as an equality for small x).

Note: This gives a lower bound on the Ramsey number $R(3, \ell) \gtrsim \ell^2 / \log^2(\ell)$, where $R(k, \ell)$ is the smallest n so that any coloring of K_n must contain either a red k -clique or

blue ℓ -clique.

Hint: Consider coloring the vertices red with probability p and blue with probability $1 - p$, where $p = \Theta(1/\sqrt{n})$.

Hint: You may run into a situation where you want to count the number of sets over vertices of size ℓ that intersect a set T of size 3 in at least two places. Try trivially bounding this by $\binom{n}{\ell}$, which is just the number of sets of size ℓ .

Hint: If you would like an oracle to tell you how you should choose the $x(A)$, ask Mary during class...

Group Work: Solutions

Following the hint, we take $p = C/\sqrt{n}$ for some constant C . For a set T of three vertices, let A_T be the event that the triangle defined by T is all red. For a set S of ℓ vertices, let B_S be the event that the K_ℓ is all blue. Then

$$\Pr[A_T] = p^3 \quad \Pr[B_S] = (1 - p)^{\binom{\ell}{2}}$$

Next we consider the dependency graph. We get a legit dependency graph if we put an edge:

- between A_T and $A_{T'}$ whenever T, T' share an edge (for each T , at most $3n$ such T');
- between A_T and B_S so that T and S share an edge (for each T , there are trivially at most $\binom{n}{\ell}$ such S ; for each S , there are at most $\binom{\ell}{2}n$ such T);
- between B_S and $B_{S'}$ so that S, S' share an edge (for each S , there are trivially at most $\binom{n}{\ell}$ such S').

Now, we need to assign values $x(A_T)$ and $x(B_S)$. Let's let $x(A_T) = x$ for all T and $x(B_S) = y$ for all S . We need to figure out what x and y should be. To actually solve the problem, we'd work backwards, but with the advantage of hindsight we'll pick

$$x = C'n^{-3/2} \quad y = 1/\binom{n}{\ell},$$

for some constant C' . Now we try to apply the Asymmetric LLL. We want:

$$p^3 \leq x \cdot (1 - x)^{3n} (1 - y)^{\binom{n}{\ell}} \tag{1}$$

and

$$(1 - p)^{\binom{\ell}{2}} \leq y \cdot (1 - x)^{n\binom{\ell}{2}} (1 - y)^{\binom{n}{\ell}} \tag{2}$$

Let's start with the first one. By our choice of y , we have $(1 - y)^{\binom{n}{\ell}} \approx 1/e$. By our choice of x ,

$$(1 - x)^{3n} = (1 - C'n^{-3/2})^{3n} \approx e^{-1/\sqrt{n}} \approx 1 - C'/\sqrt{n}$$

for large enough n .

Plugging in our choice of p and x , we get that to have (1), it's enough to have

$$C^3 n^{-3/2} \leq \frac{1}{e} \cdot n^{-3/2} \cdot (1 - C'/\sqrt{n})$$

This will be true (for large enough n) as long as we choose C small enough. (eg, $C = 1/e$ will do it comfortably).

Now we turn our attention to (2). Again we have $(1 - y)^{\binom{n}{\ell}} \approx 1/e$. Using our choice of x , we have

$$(1 - x)^{n \binom{\ell}{2}} \approx \exp(-C' \binom{\ell}{2} / \sqrt{n}).$$

Using our choice of p , we have

$$(1 - p)^{\binom{\ell}{2}} \approx \exp(-C \binom{\ell}{2} / \sqrt{n})$$

Thus, (2) reads

$$\exp(-C \binom{\ell}{2} / \sqrt{n}) \leq \frac{1}{\binom{n}{\ell}} \cdot \exp(-C' \binom{\ell}{2} / \sqrt{n}) \cdot (1/e).$$

Taking logs of both sides, we need

$$\frac{C \binom{\ell}{2}}{\sqrt{n}} \geq \ell \log(n) + \frac{C' \binom{\ell}{2}}{\sqrt{n}} + 1$$

Rearranging, we need that

$$(C - C') \binom{\ell}{2} \geq \sqrt{n}(\ell \log(n) + 1),$$

for which it suffices to have

$$(C - C')\ell \geq 2\sqrt{n}(\log(n) + 1/\ell).$$

Thus, by choosing C' small enough (say, $C' = C/2$), it's enough to have

$$\ell \geq C''(\sqrt{n} \log n)$$

for some third constant C'' . Since we assumed that $\ell \gg \sqrt{n} \log n$, this is true.