

Q: What is the class NP?

Defn: defined for decision problems

e.g. is given a graph, is it 3-COLOURABLE?

Here $\mathcal{L} = \{ G : G \text{ is 3-colourable} \}$

[NP is a class of languages which admits a poly-time non-deterministic Turing machine.]

In contrast P is the class of languages identifiable with deterministic, poly-time Turing machines

Equivalent viewpoint of the class NP is a "prover-verifier language"


NP is the class of languages L st \exists deterministic polynomial time Algo (VERIFIER)

st

$\forall x \in L, \exists$ proof $\pi(x)$ of poly(|x|) bits such that

$$V(x, \pi(x)) = 1$$

$\forall x \notin L$, no proof should make V accept
 (ies $\forall \pi(x)$, $V(x, \pi(x)) = 0$)



Given graph, is it 3-colourable?

2COL: Given a graph, is it 2-colourable?
 $L = \{ \text{2-colourable graphs} \}$

Is $P = NP$

Most longstanding problems
 in complexity theory.


{ Widely believed that $P \neq NP$. }

Cook-Levin Theorem [1971]

3SAT \equiv NP-Complete.

Given a CNF where each clause is
 an OR of 3 literals (x_i or \bar{x}_i)
 check if it is satisfiable or not.

When is a problem NP-complete?



A problem Q is NP Complete if

a) it is in NP

b) every other problem in NP can be reduced to Q in Poly-time.

Meaning,

if we discover a poly-time algo for SAT, then we can obtain poly-time algs for any problem in NP by using Cook-Levin Thm.

This notion of NP completeness gave a way of showing how "hard" some decision problems are.

3COL is NP-complete
How would we show something like this?

a) 3COL \in NP [Proof is the coloring]

b) Every other problem in NP can reduce to 3COL in polytime?

(b) seems tricky, but is not due to Cook-Levin
Simply reduce 3SAT to 3COL in polytime

(1c)

Given I of 3SAT, output a graph $G(I)$
in poly-time such that

I is satisfiable $\iff G$ is 3-colorable.

Cheer! We have a way of characterizing
"hardness" of decision problems by
showing NP-completeness.

For longest time, it was not clear how
to extend these ideas / notions to
optimization / approximation-type problems. ☹️

Ex:

We know 3-SAT is NP-complete.

Let's look at Max-3SAT:

Given an instance, find the

0... assignment which satisfies
Max # clauses. ?

- Can't admit poly-time algo (if $P \neq NP$)
Can it admit a PTAS?

(i.e.) \forall constant $\epsilon > 0$, can we get a $(1-\epsilon)$ -
approximation runs in poly-time?

Similarly, can ask for
Graph Colouring

Given graph G , colour it with
fewest colours

What's the best approximation algo for
this? Is there a
 $(1+\epsilon)$ -approximation? Is there
a 2-approximation?

Another side of the story.

{ How to analyze the class NP from
the perspective of the verifier?

How far can we restrict the verifier and
still admit the NP-class?

(long story for the motivation)

One attempt [ALMSS '98]

Probabilistically Checkable Proof PCP

$PCP[r(n), s(n)]$ is the class of all languages L s.t. \exists poly-time verifier V

$x \in L \Rightarrow \exists$ proof $\pi(x)$ of size $\text{poly}(|x|)$
such that V probes only $s(|x|)$ bits of the proof and accepts with prob. 1.

$x \notin L \Rightarrow \nexists$ proofs, V rejects after checking $s(|x|)$ bits of the proof with prob. $\frac{1}{2}$.

The verifier V is restricted to probe only $s(n)$ bits of the proof.

Hence we give it some power to toss $r(n)$ random coins.

So V can look at x and the
 $r(n)$ coins to determine which
 $s(n)$ bits to probe.

PCP $[r(n), s(n)]$
poly time V verifier which tosses $r(n)$ coins
randomized

$x \in L \Rightarrow \exists \text{ proof } \pi(x)$
 $\downarrow V$ probes $s(n)$ bits of proof π and
accepts w.p. 1

$x \notin L \Rightarrow \nexists \text{ proof } \pi, V$ rejects w.p. $\geq 1/2$.

REMARKABLE THM [ALMSS '98]

$NP = PCP [O(\log n), O(1)] !!!$

\Rightarrow can check if $x \in L$ by looking
at only constant bits of the proof!

Has beautiful connections to hardness of
approximations

Opened the door to studying if
3SAT, COL have PTAS, etc.

Goal: Hardness of Approximation of Max3SAT

Prior to PCPs, was not known if

{ Max 3SAT admits a $(1-\epsilon)$ -approximation
for any constant $\epsilon > 0$.

(i.e.) PTAS?

Emergence of PCPs \Rightarrow PTAS NOT POSSIBLE

There is some fixed $\epsilon > 0$ for which it
is NP-hard to approximate Max3SAT
to a factor better than $(1-\epsilon)$.

Subsequent improvements to PCP Machinery

$\forall \epsilon > 0$, it is NP-hard to get a

$(\frac{7}{8} + \epsilon)$ -approx to Max3SAT.

And we know trivial $\frac{7}{8}$ -approximation.

Q: What does it mean? ^{poly-time}

A: If you design a $(\frac{7}{8} + \epsilon)$ -approx for $\epsilon > 0$
then we can design a poly-time
algo for 3SAT.

PCP view of the NP class allows us
to go from such approx-hardness
to classical NP-completeness.

Q: How are we going to show APX-hardness
of Max 3SAT?

Idea: create an intermediate decision problem
and show it is NP-complete.

Problem

Gap 3SAT_{C,S}:

Given an instance I of 3SAT, output

① Yes if there is an assignment
satisfying $\geq C$ fraction of clauses

② No if no assignment can satisfy
 $\geq S$ fraction of clauses

③ YES/NO if intermediate values.

Ex: for $C = 1$, $S = 0.9$

Gap 3SAT_{1,0.9}:



Intuitively easier than regular 3SAT,
where we need to give

Correct answer in all regimes of satisfiability.

Using PCP Thm, we can show that
Gap 3SAT_{1, 1-ε} is NP-complete for
some constant $\epsilon > 0$.

Using \star , easy to show $(1-\epsilon)$ -hardness
of Max 3SAT.

SpS \exists A which is a $(1-\epsilon)$ -approx for
Max 3SAT

Given I of Gap 3SAT_{1, 1-ε}, run A on I .
if it finds an assignment satisfying
 $\geq (1-\epsilon)$ fraction of clauses,
Output YES
- Else output NO


Todo
Show that Gap 3SAT_{1, 1-ε} is NP-complete

Proof

From PCP Theorem + reduction

We know that 3COL is NP-complete.

We'll reduce it to $\text{Gap } 3\text{SAT}_{1,1-\epsilon}$.
 Traditional reductions don't give such "gaps"
 PCP view comes to our rescue.



we'll use the PCP view of 3-coloring.

\exists verifier V which tosses $O(\log n)$ -random coins, and looks at $O(1)$ bits of the proof to decide if $G \in 3\text{COL}$.

- if $G \in 3\text{COL}$, \exists proof st $\Pr_{\text{coins}} [V \text{ accepts}] = 1$
- if $G \notin 3\text{COL}$, \forall proofs $\Pr_{\text{coins}} [V \text{ accepts}] \leq 1/2$.

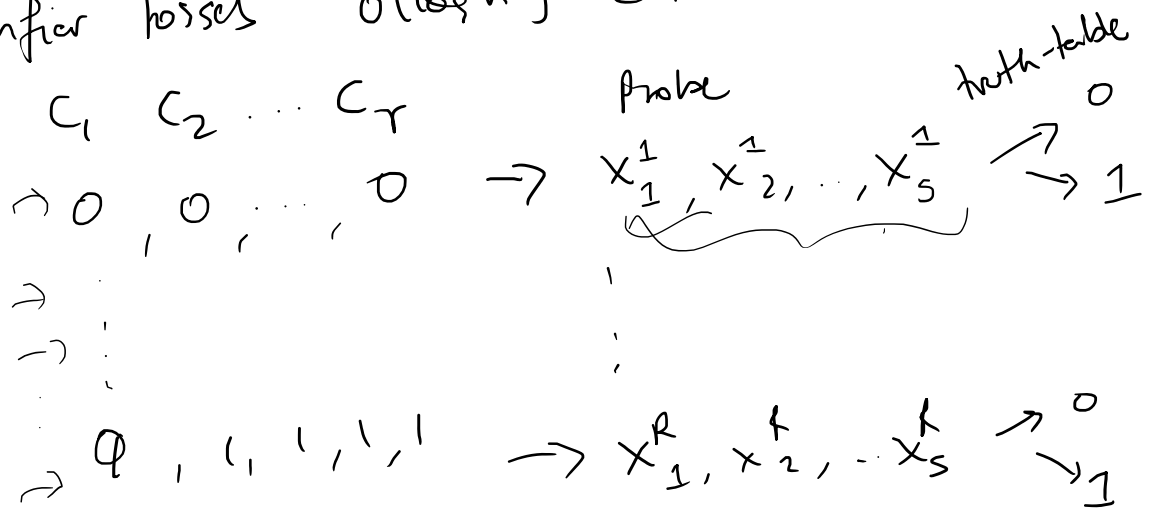
Our reduction will create an instance of $\text{Gap } 3\text{SAT}_{1,1-\epsilon}$ as follows :-

Imagine that proof provided is of the form

$$x_1, x_2, x_3, \dots, x_N$$

where $N = \text{poly}(n)$. proof size

Verifier losses $O(\log n)$ - Coins



We'll form a 3CNF from these truth tables.

We'll convert each truth table to a collection of clauses. (small #)

For example

x_1	x_{10}	x_{15}	x_{20}	output of V
0	1	1	1	0
1	0	1	0	0
1	1	1	1	0

otherwise accept.

Verifier

is checking

$$(\bar{x}_1 \wedge x_{10} \wedge x_{15} \wedge x_{20}) \vee \dots$$

$$\text{reject} \Leftrightarrow (x_1 \wedge \bar{x}_0 \wedge x_{15} \wedge \bar{x}_{20}) \vee (x_1 \wedge x_{10} \wedge x_{15} \wedge x_{20})$$

$$\text{accept} \Leftrightarrow (x_1 \vee \bar{x}_0 \vee \bar{x}_{15} \vee \bar{x}_{20}) \wedge (\bar{x}_1 \vee x_{10} \vee \bar{x}_{15} \vee x_{20}) \wedge (\bar{x}_1 \vee \bar{x}_{10} \vee \bar{x}_{15} \vee \bar{x}_{20})$$

[Can use some small # auxiliary variables to make it a 3CNF.]

In total, we'll generate $\text{poly}(n)$ $\text{poly}(s)$ clauses but s is constant.

This is the desired Gap 3SAT instance \tilde{I} .

Final part of proof:

$G \in 3\text{COL}$, then all clauses of \tilde{I} are satisfied

$G \notin 3\text{COL}$, then $\leq (1-\epsilon)$ -fraction of clauses of \tilde{I} are satisfied

PCP's are intimately tied to hardness of Approx

How to show hardness of Apix?

For optimization problem P (let's say maximization)
 Create a Gap $P_{c,s}$ problem (decision problem)
 & show it is NP-complete

\Rightarrow NP-hardness of factor $\left(\frac{s}{c}\right)$ for problem P .

Try to find c, s such that $\frac{s}{c}$ is as small as possible

This gap is related to accept prob. of verifier in PCP theorem.

"Parallel Repetition Theorem"

lots of research to trade-off $r(n)$, $s(n)$, accept probability
 culminated in a very neat abstraction called label cover problems.

equivalent form of PCP Thm + Parallel Repetition where verifier probes only 2 locations of the proof.


[Caveat: proof is no longer bit string, its over a larger alphabet]

Max Label Cover Problem:

Given a bipartite graph $G = [U, V, E]$
 and a "projection function" $f_{u \rightarrow v} : [L] \rightarrow [K]$
 on each edge $(u, v) \in E$,

goal

Pick labels
 for each vertex (ie)
 $l(u) \in [L]$ for all $u \in U$
 $l(v) \in [K]$ for all $v \in V$




Max # edges for which
 labels are aligned wrt f .

$$(ie) \quad f_{u \rightarrow v}(l(u)) = l(v)$$

Think of K and L as being constants

Graph Label Cover, η

Given an instance of label cover,
 can we decide if \exists labelling
 which can satisfy all edges vs
 if all labellings can satisfy $\leq \eta$ fraction
 of edges?




THEOREM:

For any constant $\eta > 0$, \exists constant
 k, L bounded by $\text{poly}(1/\eta)$

k, L bounded by $\text{poly}(1/\eta)$
 Such that the $\text{GapLabelCover}_{1,\eta}$ is
 NP-complete.

Moreover,
 this holds
 even
 when
 $|V|=|E|$
 and
 graph is
 regular

Many, if not most, hardness of approx.
 results come with the GapLabelCover
 as starting point.



Example Application (Hardness of Max-Coverage)

Problem:

Given a set system (U, \mathcal{F}) where
 U is universe of elements and
 \mathcal{F} is collection of subsets of U ,
 and given parameter k , goal:
 choose k sets from \mathcal{F} , say

$$S_1, S_2, \dots, S_k \in \mathcal{F} \quad \text{to maximize}$$

$$\left| \bigcup_{i=1}^k S_i \right| \quad (\text{ie the \# elts covered by them}).$$

Algorithmic Ideas :-

① greedy algo: for $t=1, 2, \dots, k$
 choose set which covers
 max # uncovered elts

② \lfloor -rounding?

③ local search: Start with k random sets, and swap in a new set for an existing set if coverage improves

Thm
Greedy Algo is a $(1 - \frac{1}{e})$ -approximation

Thm
For any fixed $\epsilon > 0$, it is NP-complete to design a poly-time $(1 - \frac{1}{e} + \epsilon)$ -approx.

↗ PCP + label cover viewpoint
↘

Thm 2
For any fixed $\epsilon > 0$, it is NP-complete to design a poly-time $(\ln n - \epsilon)$ -approx for Set cover. (min # sets to cover all elems)

Recall: greedy algo is $\ln n$ -approx.

Next lecture

Briefly outline the redn. from Graph Label Cover \rightarrow Gap Max Coverage.
to show slightly worse factors of $(\frac{3}{4} + \epsilon)$ hardness of Approx.

Given an instance $\mathcal{I} = \{G = (V, E), \{f_{uv}\}\}$ of Graph Label Cover problem, we'll create an instance $\mathcal{I}' = \{X, \mathcal{F}, k\}$ of Gap Coverage st.

$$\textcircled{1} \begin{cases} \text{if } \text{Opt}(I) = 1 & \implies \text{Opt}(I') = 1 \\ \text{Opt}(I') \leq \gamma & \implies \text{Opt}(I) \leq \frac{3}{4} + \epsilon. \end{cases}$$

Here
 Opt = fraction of
 satisfied edges

Here
 Opt = fraction of
 covered elements.


From this,

NP-completeness of Gap Label Cover, we can get
 NP-completeness of Gap Coverage.

$\implies \left(\frac{3}{4} + \epsilon\right)$ - hardness of approx for MaxCover.

Todo

Come up with I' and prove $\textcircled{*}$ for it.



From I , we want I' such that

there is some correspondence between
 assigning a label l to u and
 picking a set in \mathcal{I}' .

We'll create \mathcal{I}' such that

there is a set $S_{u,\alpha}$ for all $u \in V$
 $\alpha \in [L]$

and similarly

$S_{v,\beta}$ for all $v \in V$
 $\beta \in [K]$

In total

$$\begin{aligned} \# \text{ sets} &= |V| \cdot L + |V| \cdot K \\ &= |V| (L + K) \\ &= n(L + K) \end{aligned} \left. \vphantom{\begin{aligned} \# \text{ sets} &= |V| \cdot L + |V| \cdot K \\ &= |V| (L + K) \\ &= n(L + K) \end{aligned}} \right\} \begin{array}{l} \because |U| = |V| = n \\ \text{in input} \\ \text{graph.} \end{array}$$

Next, we'll create some sort of association
 between edges in \mathcal{I} with elements
 of \mathcal{I}' .

For each edge $(u,v) \in G$, we'll create
 a number of elements.

2^k elements corresponding to k -bit
 strings

We'll refer to these elements as

$e_x^{(u,v)}$ where (u,v) is edge
 x is a k -bit

How many elements in \mathcal{I}' have we created? string.

$$\# \text{edges} \cdot 2^k = d \cdot n \cdot 2^k$$

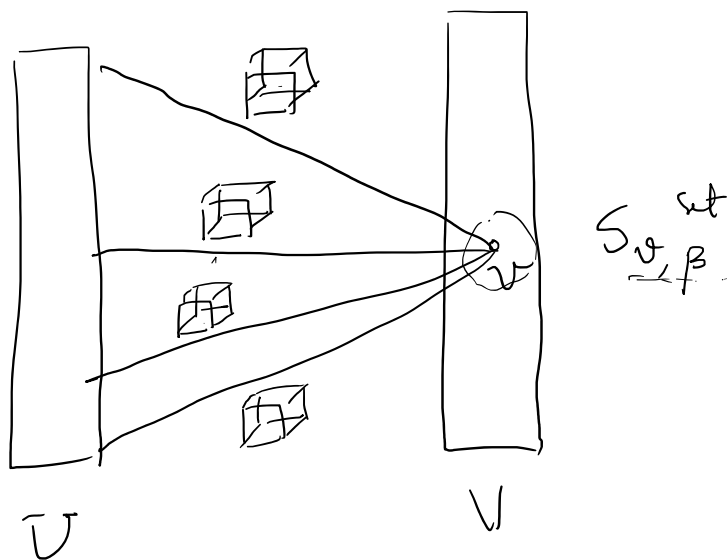
\therefore graph is d -regular.

We've defined X (universe of elts) in \mathcal{I}'
 " " \mathcal{F} (coll. of sets).

last: When does a set $S_{u,\alpha}$ cover an element?

Why, when does a set $S_{v,\beta}$ cover an element?

let's start with $S_{v,\beta}$.



RULE 1: Firstly $S_{v,\beta}$ can only cover the elements on the

edges incident to v .

$$\Rightarrow |S_{v, \beta}| \leq d \cdot 2^k \quad (\text{trivial bound})$$

Moreover,

consider edge (u, v) and all
elts $e_x^{(u,v)}$ for $x \in \{0,1\}^k$

RULE 2

let's make $S_{v, \beta}$ cover all elts where
the β th bit of $x = 1$.

In particular

$$e_x^{(u,v)} \in S_{v, \beta} \iff x[\beta] = 1 \quad (\beta\text{th bit of } x).$$

$S_{u, \beta}$ is simply all such elts

$$\Rightarrow |S_{u, \beta}| = d \cdot 2^{k-1}$$

Similarly, sets for all $v \in V$,
all $1 \leq \beta \leq k$.

Next we'll define sets for U


consider $S_{u, \alpha}$ for $u \in U$, $\alpha \in [L]$.

$$|S_{u, \alpha}| = 0$$

$$C_x^{(u,v)} \in S_{u,v} \text{ iff}$$

$$x[f_{u,v}(x)] = 0$$

(ie) the $f_{u,v}(x)^k$ bit of $x = 0$.



$$|S_{u,v}| = d \cdot 2^{k-1}$$

Also set k (in Gap k coverage) to be $\lceil 2n \rceil$.

That defines I' . Why is it useful?

CLAIM 1

I is satisfiable, I' is fully coverable

$$(10) \text{OPT}(I) = 1 \Rightarrow \text{OPT}(I') = 1$$


Proof

Let's consider the optimal labelling for I .
 It can satisfy all edges.

(10) $l^*(u)$, $l^*(v)$ are optimal labels
 $u \in U$, $v \in V$

$\forall (u, v) \in G$, it holds that

$$\boxed{f_{u,v}(l^*(u)) = l^*(v)} = \beta \quad (1 \leq \beta \leq k)$$



Using this labeling l^* , can we get a good solⁿ for I' (Gap k-Coverage)?

We can pick one set per vertex acc. l^* labeling.


$\forall u \in U$, choose $S_{u, l^*(u)}$ and

$\forall v \in V$, choose $S_{v, l^*(v)}$.

#sets selected = $2n = k$.

$S_{v, \beta}$ covers all $e_x^{(u,v)}$ st $x[\beta] = 1$

$S_{u, l^*(u)}$ covers all $e_x^{(u,v)}$ st $x[\beta] = 0$.

\Rightarrow together they cover all elts 

Remains to show:-

$$\text{Opt}(I) \leq \eta \Rightarrow \text{Opt}(I') \leq \frac{3}{4} + \epsilon.$$

instead we'll show

$$\text{Opt}(I') > \frac{3}{4} + \epsilon \Rightarrow \text{Opt}(I) > \eta.$$

Now we'll show it:-

if there is a good solⁿ for
gap coverage covering $\geq (\frac{3}{4} + \epsilon)$

fraction of elts, then we
can recover a good
labeling satisfying

$\geq \eta$ fraction of edges.

High-level sketch :-

Let S^* be a good cover $\geq (\beta + \epsilon)$ covering fraction of eds.

$$|S^*| = k = 2n.$$

Cheating assumption S^* is such that it picks one set per vertex of the graph G of I .

Then we ask : Can we get a good labeling for I .


If S^* contains $S_{u,\alpha}$, then assign label $l^*(u) = \alpha$.

Why if S^* contains $S_{v,\beta}$ then assign label $l^*(v) = \beta$.

Then claim : $\geq 4\epsilon$ fraction of edges.

we

need to have
a satisfied labeling.



Any badly labeled edge can cover
only $\leq 3\epsilon$ fraction of
edges.

But since S^* covers $\geq (\frac{3}{4} + \epsilon)$ fraction of
edges,

there must be a good # of satisfied
edges 😊

$$(1 - \delta) \cdot \frac{3}{4} + \delta \cdot 1 = \frac{3}{4} + \epsilon$$

$$\frac{3}{4} + \frac{\delta}{4} = \frac{3}{4} + \epsilon$$

$$\boxed{\delta = 4\epsilon}$$