A ZK protocol for graph 3-coloring:

contains n nodes, m edges

prover (G)

- let $k_i \in \{0,1,2\}$ be a 3-coloring of G
- choose random permutation $\tau \xleftarrow{R} \mathrm{Perm}[\{0,1,2\}]$
- for $i \in [n]$: $(c_i, \pi_i) \leftarrow \mathrm{Commit}(\sigma, \tau(k_i))$ for random r_i
- reject if $(i,j) \notin E$

verifier (G)

$\sigma \leftarrow \mathrm{Setup}(1^\lambda)$

$(i,j) \xleftarrow{R} E$

Messages:

verifier → prover: σ

prover → verifier: $c_1, \dots, c_n$

verifier → prover: i, j

prover → verifier: $(\tau(k_i), \pi_i), (\tau(k_j), \pi_j)$, where $\tau(k_i) = k_i'$ and $\tau(k_j) = k_j'$

accept if $k_i' \neq k_j'$ and $k_i', k_j' \in \{0,1,2\}$

$\mathrm{Verify}(\sigma, c_i, k_i', \pi_i) = 1 = \mathrm{Verify}(\sigma, c_j, k_j', \pi_j)$

reject otherwise

Intuitively: Prover commits to a coloring of the graph
Verifier challenges prover to reveal coloring of a single edge
Prover reveals the coloring on the chosen edge and opens the entries in the commitment

Completeness: By inspection [if coloring is valid, prover can always answer the challenge correctly]

Soundness: Suppose G is not 3-colorable. Let $k_1, \dots, k_n$ be the coloring the prover committed to. If the commitment scheme is statistically binding, $c_1, \dots, c_n$ uniquely determine $k_1, \dots, k_n$ (except with prob. $1 - \mathrm{negl.}$). Since G is not 3-colorable, there is an edge $(i,j) \in E$ where $k_i = k_j$ or $i \notin \{0,1,2\}$ or $j \notin \{0,1,2\}$. [Otherwise, G is 3-colorable with coloring $k_1, \dots, k_n$.] Since the verifier chooses an edge to check at random, the verifier will choose (i,j) with probability $1/|E|$. Thus, if G is not 3-colorable,

$$\Pr[\text{verifier rejects}] \geq \frac{1}{|E|}$$

Thus, this protocol provides soundness $1 - \frac{1}{|E|}$. We can repeat this protocol $O(|E|^2)$ times sequentially to reduce soundness error to

$$\Pr[\text{verifier accepts proof of false statement}] \leq \left(1 - \frac{1}{|E|}\right)^{|E|^2} \leq e^{-|E|} = e^{-m} \quad [\text{since } 1 + x \leq e^x]$$

Zero Knowledge: We need to construct a simulator that outputs a valid transcript given only the graph G as input.

Let V^* be a (possibly malicious) verifier. Construct simulator S as follows:

1. Run V^* to get σ^*.
2. Choose $\kappa_i \leftarrow \{0,1,2\}$ for all $i \in [n]$.
   Let $(c_i, \pi_i) \leftarrow \text{Commit}(\sigma^*, \kappa_i)$
   Give $(c_1, \dots, c_n)$ to V^*.
   } Simulator does not know coloring so it commits to a random one
3. V^* outputs an edge $(i,j) \in E$
4. If $\kappa_i \neq \kappa_j$, then S outputs $(\kappa_i, \kappa_j, \pi_i, \pi_j)$.
   Otherwise, restart and try again (if fails λ times, then abort)

Simulator succeeds with probability $2/3$ (over choice of $\kappa_1, \dots, \kappa_n$). Thus, simulator produces a valid transcript with prob. $1 - \frac{1}{3^\lambda} = 1 - \text{negl}(\lambda)$ after λ attempts. It suffices to show that simulated transcript is indistinguishable from a real transcript.

- Real scheme: prover opens κ_i, κ_j where $\kappa_i, \kappa_j \xleftarrow{R} \{0,1,2\}$ [since prover randomly permutes the colors]
- Simulation: κ_i and κ_j sampled uniformly from $\{0,1,2\}$ and conditioned on $\kappa_i \neq \kappa_j$, distributions are identical

In addition, (i,j) output by V^* in the simulation is distributed correctly since commitment scheme is computationally-hiding (e.g. V^* behaves essentially the same given commitments to a random coloring as it does given commitment to a valid coloring

If we repeat this protocol (for soundness amplification), simulator simulate one transcript at a time

Summary: Every language in NP has a zero-knowledge proof (assuming existence of OWFs)

Bit commitments from PRG. Let $G: \{0,1\}^\lambda \to \{0,1\}^{3\lambda}$ be a PRG.

Setup(1^λ): Sample $\sigma \xleftarrow{R} \{0,1\}^{3\lambda}$

Commit(σ, m): sample $s \xleftarrow{R} \{0,1\}^\lambda$

if $m = 0$: output $G(s)$ $\quad \pi = s$

if $m = 1$: output $G(s) \oplus \sigma$ $\quad \pi = s$

Verify(σ, m, c, π): if $m = 0$, check that $G(\pi) = c$

if $m = 1$, check that $G(\pi) \oplus \sigma = c$

Hiding follows by PRG security.

Binding follows by union bound:

$$\Pr_{\sigma \xleftarrow{R} \{0,1\}^{3\lambda}} \left[\exists\, s_0, s_1 \in \{0,1\}^\lambda : G(s_0) = G(s_1) \oplus \sigma \right] = \frac{2^{2\lambda}}{2^{3\lambda}} = \frac{1}{2^\lambda}$$