Requirements:

- <u>Correctness</u>: for all messages m:

$$\Pr[\sigma \leftarrow \mathrm{Setup},\ (c,\pi) \leftarrow \mathrm{Commit}(\sigma, m);\ \mathrm{Verify}(\sigma, c, m, \pi) = 1] = 1$$

- <u>Hiding</u>: for all common reference strings $\sigma \in \{0,1\}^n$ and all efficient A, following distributions are computationally indistinguishable:

adversary		challenger ($b \in \{0,1\}$)
	$\xleftarrow{\ \sigma\ }$
	$\xrightarrow{\ m_0, m_1\ }$
		$(c,\pi) \leftarrow \mathrm{Commit}(\sigma, m_b)$
	$\xleftarrow{\ c\ }$
$b' \in \{0,1\}$	

$$\left|\Pr[b' = 1 \mid b = 0] - \Pr[b' = 1 \mid b = 1]\right| = \mathrm{negl}(\lambda)$$

- <u>Binding</u>: for <u>all</u> adversaries A, if $\sigma \leftarrow \mathrm{Setup}$, then

$$\Pr[(m_0, m_1, c, \pi_0, \pi_1) \leftarrow A :\ m_0 \neq m_1 \text{ and } \mathrm{Verify}(\sigma, c, m_0, \pi_0) = 1 = \mathrm{Verify}(\sigma, c, m_1, \pi_1)] = \mathrm{negl}$$

A ZK protocol for graph 3-coloring:

prover (G) — contains n nodes, m edges		verifier (G)
- let $K_i \in \{0,1,2\}$ be a 3-coloring of G	$\longleftarrow$	$\sigma \leftarrow \mathrm{Setup}(1^\lambda)$
- 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	$\xrightarrow{\ c_1, \dots, c_n\ }$	$(i,j) \xleftarrow{R} E$
- reject if $(i,j) \notin E$	$\xleftarrow{\ i,j\ }$
	$\xrightarrow{\ (\tau(K_i), \pi_i),\ (\tau(K_j), \pi_j)\ }$ with $\tau(K_i) = K_i'$, $\tau(K_j) = K_j'$

$\hookrightarrow$ 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. $1 - \text{negl.}$). 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 $K_i \leftarrow \{0,1,2\}$ for all $i \in [n]$.
   Let $(c_i, \pi_i) \leftarrow \text{Commit}(\sigma^*, K_i)$
   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 $K_i \neq K_j$, then S outputs (K_i, K_j, π_i, π_j).
   Otherwise, restart and try again (if fails λ times, then abort)

Simulator succeeds with probability $2/3$ (over choice of $K_1, \dots, K_n$). 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 K_i, K_j where $K_i, K_j \xleftarrow{R} \{0,1,2\}$ [since prover randomly permutes the colors]
- Simulation: K_i and K_j sampled uniformly from $\{0,1,2\}$ and conditioned on $K_i \neq K_j$, 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 PRGs)

PRGs imply commitments