

Zero-Knowledge Proofs for Classical Planning Problems

Augusto B. Corrêa, Clemens Büchner, Remo Christen

University of Basel, Switzerland

AAAI 2023

propositional classical planning tasks:

- propositional variables: $\{x, y\}$
- initial state: $\{x, \neg y\}$
- goal: $\{y\}$
- actions:
 - $a_1: \quad pre = \{\}$
 $\quad \quad \quad eff = \{\neg x\}$
 - $a_2: \quad pre = \{\neg x\}$
 $\quad \quad \quad eff = \{y\}$
- plan: $\pi = \langle (s_0, a_1, s_1), (s_1, a_2, s_2) \rangle$

Planning without Revealing Solutions

I know a plan for some planning task.

But I won't show you.

Can I convince you that I really know a plan?

How do I do that without revealing anything about the plan?

Planning without Revealing Solutions

I know a plan for some planning task.

But I won't show you.

Can I convince you that I really know a plan?

How do I do that without revealing anything about the plan?

Zero-Knowledge Proofs!

Prover:

- claims to have a plan π (somehow)
- wants to prove this without revealing anything about π

Verifier:

- wants to check that **Prover** is not lying
- do it efficiently

Prover:

- claims to have a plan π (somehow)
- wants to prove this without revealing anything about π

Verifier:

- wants to check that **Prover** is not lying
- do it efficiently

bounded plan existence:

Given a planning task Π and $k \in \mathbb{N}$ such that $k \leq \text{poly}(|\Pi|)$,
is there a plan π for Π with $|\pi| \leq k$?

Prover:

- claims to have a plan π (somehow)
- wants to prove this without revealing anything about π

Verifier:

- wants to check that **Prover** is not lying
- do it efficiently

bounded plan existence:

Given a planning task Π and $k \in \mathbb{N}$ such that $k \leq \text{poly}(|\Pi|)$, is there a plan π for Π with $|\pi| \leq k$?

- no limitation on k if you reduce to QBF protocols
- **Prover** needs to stronger than in our case; algebraic

completeness: **Verifier** will never reject plan of an honest **Prover**

soundness: **Verifier** might be fooled with low probability

zero-knowledge: **Verifier** learns nothing about the plan

efficiency: protocol executed efficiently by **Verifier** and **Prover**

overview of all steps:

- **Prover** obfuscates Π
- **Verifier** chooses between
 - verifying if obfuscation was done correctly
 - verifying one transition of the claimed plan π
- if the chosen verification succeeds, **Verifier** accepts π otherwise, rejects

Prover communicates via encrypted messages

- opened using some key

- propositional variables: $\{x, y\}$
- initial state: $\{x, \neg y\}$
- goal: $\{y\}$
- actions:
 - $a_1: \quad pre = \{\}$
 $\quad \quad \quad eff = \{\neg x\}$
 - $a_2: \quad pre = \{\neg x\}$
 $\quad \quad \quad eff = \{y\}$
- plan: $\pi = \langle (s_0, a_1, s_1), (s_1, a_2, s_2) \rangle$

Prover Obfuscates Π

transform Π into $\hat{\Pi}$ such that it is hard to map $\hat{\Pi}$ back to Π

variables	initial state	goal	actions	
$\{x, y\}$	$\{x, \neg y\}$	$\{y\}$	$\{\} \rightarrow \{\neg x\}$	$\{\neg x\} \rightarrow \{y\}$

Prover Obfuscates Π

transform Π into $\hat{\Pi}$ such that it is hard to map $\hat{\Pi}$ back to Π

variables	initial state	goal	actions	
$\{x, y\}$	$\{x, \neg y\}$	$\{y\}$	$\{\} \rightarrow \{\neg x\}$	$\{\neg x\} \rightarrow \{y\}$

make actions indistinguishable:

$\{x, y, z\}$	$\{x, \neg y, \neg z\}$	$\{y\}$	$\{\neg z\} \rightarrow \{\neg x\}$	$\{\neg x\} \rightarrow \{y\}$
---------------	-------------------------	---------	-------------------------------------	--------------------------------

transform Π into $\hat{\Pi}$ such that it is **hard to map $\hat{\Pi}$ back to Π**

variables	initial state	goal	actions	
$\{x, y\}$	$\{x, \neg y\}$	$\{y\}$	$\{\} \rightarrow \{\neg x\}$	$\{\neg x\} \rightarrow \{y\}$

make actions indistinguishable:

$\{x, y, z\}$	$\{x, \neg y, \neg z\}$	$\{y\}$	$\{\neg z\} \rightarrow \{\neg x\}$	$\{\neg x\} \rightarrow \{y\}$
---------------	-------------------------	---------	-------------------------------------	--------------------------------

change variable labels and sign:

$\{\chi, v, \zeta\}$	$\{\chi, v, \neg \zeta\}$	$\{\neg v\}$	$\{\neg \zeta\} \rightarrow \{\neg \chi\}$	$\{\neg \chi\} \rightarrow \{\neg v\}$
----------------------	---------------------------	--------------	--	--

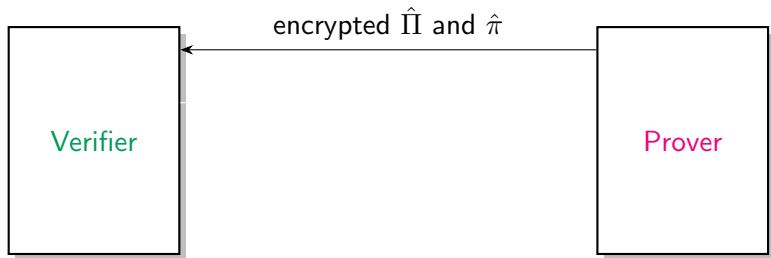
transform Π into $\hat{\Pi}$ such that it is **hard to map $\hat{\Pi}$ back to Π**

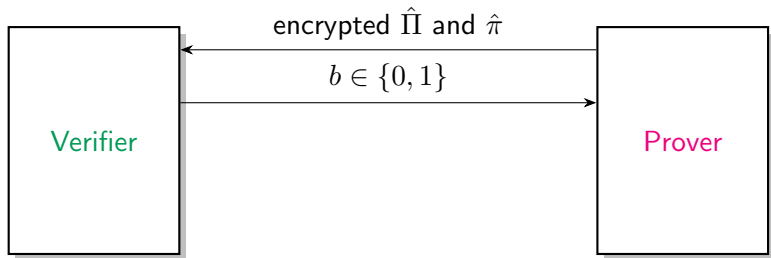
variables	initial state	goal	actions	
$\{x, y\}$	$\{x, \neg y\}$	$\{y\}$	$\{\} \rightarrow \{\neg x\}$	$\{\neg x\} \rightarrow \{y\}$
make actions indistinguishable:				
$\{x, y, z\}$	$\{x, \neg y, \neg z\}$	$\{y\}$	$\{\neg z\} \rightarrow \{\neg x\}$	$\{\neg x\} \rightarrow \{y\}$
change variable labels and sign:				
$\{\chi, \nu, \zeta\}$	$\{\chi, \nu, \neg \zeta\}$	$\{\neg \nu\}$	$\{\neg \zeta\} \rightarrow \{\neg \chi\}$	$\{\neg \chi\} \rightarrow \{\neg \nu\}$
add canonical initial and goal states: (omitted)				

transform Π into $\hat{\Pi}$ such that it is **hard to map $\hat{\Pi}$ back to Π**

variables	initial state	goal	actions	
$\{x, y\}$	$\{x, \neg y\}$	$\{y\}$	$\{\} \rightarrow \{\neg x\}$	$\{\neg x\} \rightarrow \{y\}$
make actions indistinguishable:				
$\{x, y, z\}$	$\{x, \neg y, \neg z\}$	$\{y\}$	$\{\neg z\} \rightarrow \{\neg x\}$	$\{\neg x\} \rightarrow \{y\}$
change variable labels and sign:				
$\{\chi, v, \zeta\}$	$\{\chi, v, \neg \zeta\}$	$\{\neg v\}$	$\{\neg \zeta\} \rightarrow \{\neg \chi\}$	$\{\neg \chi\} \rightarrow \{\neg v\}$
add canonical initial and goal states: (omitted)				

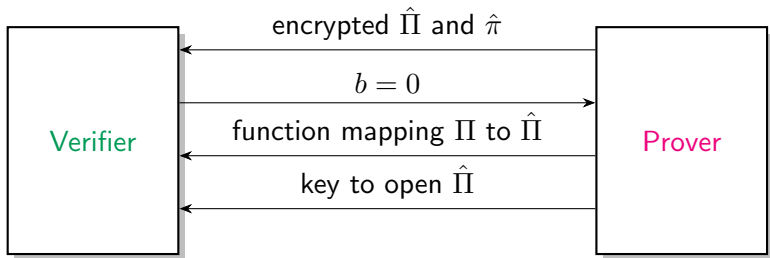
- change plan π into $\hat{\pi}$
- done in **polynomial time** (details in the paper)





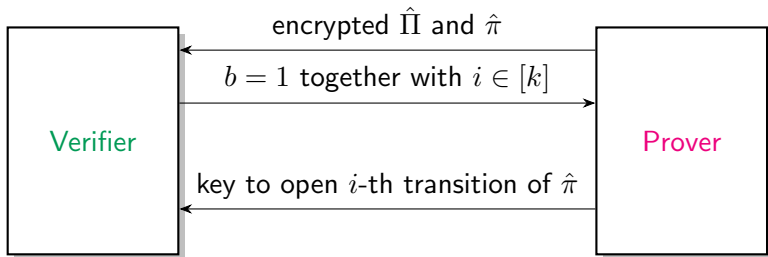
Verifier chooses a random value $b \in \{0, 1\}$

- if $b = 0$: verifier checks the transformation $\Pi \rightarrow \hat{\Pi}$
- if $b = 1$: verifier checks one randomly chosen transition of $\hat{\pi}$



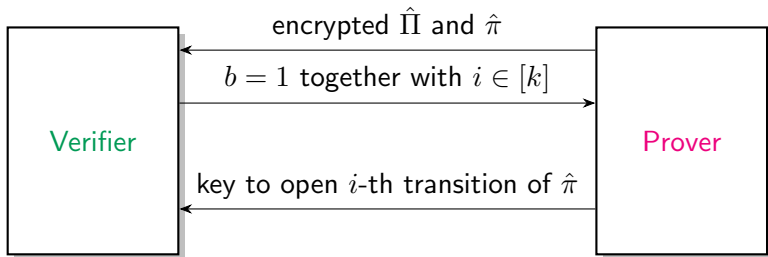
If sent function indeed maps Π into $\hat{\Pi}$, **Verifier accepts**.
Otherwise, **rejects**.

Protocol: $b = 1$



If transition (s_{i-1}, a_i, s_i) is valid, **Verifier accepts**.
Otherwise, **rejects**.

Protocol: $b = 1$



If transition (s_{i-1}, a_i, s_i) is valid, **Verifier accepts**.
Otherwise, **rejects**.

in practice, a little bit harder:

- check if s_0 and s_k are canonical initial and goal states
- check if a_i is an action of $\hat{\Pi}$

Properties (Revisited)

completeness: **Verifier** will never reject plan of an honest **Prover**

soundness: **Verifier** might be fooled with probability $1 - \frac{1}{2^k}$

- run protocol multiple times to reduce this probability

zero-knowledge: **Verifier** learns nothing about the plan

- see paper for details

efficiency: protocol executed efficiently by **Verifier** and **Prover**

- constant number of messages per execution
- only polynomial-time computation

zero-knowledge proofs for classical planning problems:

- how to prove that you have a plan without revealing it
- works for problems with polynomially long plans
- executed efficiently