Towards Non-interactive Witness Hiding

Benjamin Kuykendall   Mark Zhandry

Princeton University and NTT Research

TCC 2020

Question: Can we achieve non-interactive witness hiding proofs for all of NP in the standard model?

Answer: Almost. From appropriate assumptions we get

1. Witness hiding in 2 messages
2. Non-uniform witness hiding
3. Universal non-interactive proofs
4. Witness hiding vs witness encryption

Conclusion: Strong evidence that NIWH should exist, but no concrete and provably secure scheme from good assumptions

Table of Contents

Background

Outline of results

Technical details

1. Witness hiding in 2 messages
2. Non-uniform witness hiding
3. Universal non-interactive proofs
4. Witness hiding vs witness encryption

Proof system basics

Take any language $L \in \mathrm{NP}$ with verifier V_L and witness relation:

$$(x, w) \in R_L \;\Leftrightarrow\; V_L(x, w) \text{ accepts}$$

A protocol Π is an proof system for L, executed by two parties:

Prover: P gets input (x, w)

Verifier: V gets input x, either accepts or rejects

$P(x, w) \leftrightarrow V(x)$ denotes output of V at end of protocol

Properties of proof systems

Complete: $(x, w) \in R_L \;\Rightarrow\; P(x, w) \leftrightarrow V(x)$ accepts

Sound: $\forall P^* \; x \notin L \;\Rightarrow\; P^*(x) \leftrightarrow V(x)$ rejects

Efficient: P, V both ppt algorithms

Privacy notions for proof systems

ZK

zero knowledge

WH

witness hiding

WI

witness indistinguishable

Zero knowledge [GMR85]

Zero knowledge: any malicious verifier can be simulated

For any V^* ppt there exists S ppt such that $\forall (x, w) \in L$

$$\boxed{P(x, w) \;\; \begin{matrix} \longleftarrow \\ \longrightarrow \\ \longleftarrow \end{matrix} \;\; V^*(x)} \;\cong\; \boxed{S(x)}$$

Zero knowledge

- ✓ Strong notion of privacy
- ✓ With CRS: non-interactively from various assumptions
  [FLS90, CCH+19, PS19]
- ✗ Standard model: requires at least 3 messages
  [GO94, BLV03]

Witness indistinguishability [FS90]

Witness indistinguishability: malicious verifier does not know which of two witnesses is being used

For any V^* ppt and sequence of $(x, w), (x, z) \in R_L$

$$P(x, w) \overset{\longleftarrow}{\underset{\longleftarrow}{\longrightarrow}} V^*(x) \quad \cong \quad P(x, z) \overset{\longleftarrow}{\underset{\longleftarrow}{\longrightarrow}} V^*(x)$$

Witness indistinguishability

- ✓ Standard model: non-interactively from NIZK + HSG [DN00, BOV03]
- ✓ Useful in developing other protocols
- ✗ Not a meaningful privacy notion for all languages

Witness hiding [FS90]

Witness hiding: no malicious verifier can output a witness

Relative to distribution $\mathcal{D}$ over R_L: only makes sense if hard to find witnesses in the first place.

For any V^* ppt and $(x, w) \sim \mathcal{D}$

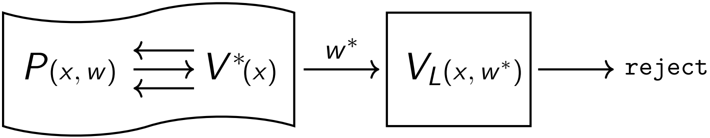

Witness hiding

- ✓ Meaningful and intuitive for any hard distribution
- ✓ With CRS: follows from NIZK
- ? Standard model: unknown

Table of Contents

Background

Outline of results

Technical details

1. Witness hiding in 2 messages
2. Non-uniform witness hiding
3. Universal non-interactive proofs
4. Witness hiding vs witness encryption

What we achieve

Four constructions that almost achieve the desired notion of NIWH

All the constructions start with a NIWI for NP and use it to construct NIWH

1. Witness hiding in 2 messages

Starting point: 2-message arguments from [Pas03]

Original security analysis: given quasipolynomially hard OWF, protocols is witness hiding when the $\mathcal{D}$-search problem is quasipolynomially hard.

New analysis: given quasipolynomially hard OWF, the protocol is witness hiding in the delayed input model when the $\mathcal{D}$-search problem is hard against non-uniform adversaries. (Result is comparable to existing work [JKKR17], but simpler construction).

2. Non-uniform witness hiding

Non-interactive proof system where prover and verifier take advice

Making a (non-standard) worst-case complexity assumption, there exists a choice of advice such that the protocol is witness hiding

But unfortunately no use in practice; unclear how to choose advice

3. Universal non-interactive proofs

Construct non-interactive proof system Π_U that is witness hiding as long as some non-interactive proof system Π' is witness hiding and provably sound

Even if Π' has an inefficient prover, Π_U is efficient

Even if Π' is non-uniform, Π_U is uniform

Unfortunately, construction above does not meet the provable soundness requirement

4. Witness hiding vs witness encryption

Non-interactive proof system for languages with unique witnesses

Either the proof system is witness hiding,
or it yields a form of witness encryption

Since witness encryption is only known from strong assumptions,
this suggests the former case is more likely

Table of Contents

Background

Outline of results

Technical details

1. Witness hiding in 2 messages
2. Non-uniform witness hiding
3. Universal non-interactive proofs
4. Witness hiding vs witness encryption

Table of Contents

Background

Outline of results

Technical details

1. Witness hiding in 2 messages
2. Non-uniform witness hiding
3. Universal non-interactive proofs
4. Witness hiding vs witness encryption

Witness hiding in 2 messages: Basic idea

To prove x with witness w output

NIWI
"$x \vee y$"
witness: w

If y is false: then proof will be sound

If y is true: then proof will be witness hiding

Witness hiding in 2 messages: Basic idea

If y is true: then proof will be witness hiding

Proof: let A be an adversary that breaks witness hiding

Let z be a witness for y. Then

$$\begin{array}{|c|}\hline \text{NIWI} \\ \hline \text{“}x \vee y\text{”} \\ \text{witness: } w \\ \hline \end{array} \cong \begin{array}{|c|}\hline \text{NIWI} \\ \hline \text{“}x \vee y\text{”} \\ \text{witness: } z \\ \hline \end{array}$$

So running A on NIWI$(x \vee y, z)$ solves the $\mathcal{D}$ search problem.

Witness hiding in 2 messages: Basic idea

Of course, y cannot be both true and false

Resolution: sample y that is true, but finding a witness is hard

To do this, we use a one-way function f and let

$$y := \exists r' \, : \, b = f(r')$$

Witness hiding in 2 messages: construction

Verifier: sample $r \sim \{0,1\}^k$ and output $b = f(r)$

Prover: output a commitment c to w along with

NIWI
"$\exists w' \; : \; (c = \mathrm{Comm}(w')) \wedge ((x, w') \in R_L \vee b = f(w'))$"   witness: w

Verifier: verify the NIWI is a valid proof of the desired statement

Witness hiding in 2 messages: complexity leveraging

Two things to prove:

Soundness: break the commitment, yielding a OWF pre-image

Witness hiding: invert the OWF and use r to generate the NIWI

Both of these adversaries are inefficient: thus witness hiding is only achieved when the OWF and commitment have carefully chosen concrete security parameters and $\mathcal{D}$ is secure against quasipolynomial time adversaries.

Witness hiding in 2 messages: delayed input model

Would prefer to use standard hardness of $\mathcal{D}$

Delayed input model: x is revealed to the verifier at the end [JKKR17]

To prove witness hiding, we can non-uniformly fix a choice of r

Witness hiding in 2 messages: removing interaction?

Note r is never used in the protocol

Thus if f is a permutation, we directly sample $b \sim \{0,1\}^k$

Gives straightforward heuristic to remove interaction:
take hash function H and run with $b = H(x)$

Can be shown secure in (non-programmable) random oracle model,
but not clear we can do any better

Table of Contents

Background

Outline of results

Technical details

1. Witness hiding in 2 messages
2. Non-uniform witness hiding
3. Universal non-interactive proofs
4. Witness hiding vs witness encryption

Non-uniform: construction

Again the verifier simply sends a NIWI of $x \vee y$

But now we fix y non-uniformly:
take it as an advice string for both prover and verifier

We fix y to be false for soundness

Non-uniform: sketching witness hiding

Fix y and take a successful adversary A_y against witness hiding

We know that the protocol is witness hiding if y were true

Thus A_y is a "proof" that y must be false

But if we believe coNP $\not\subset$ NP such "proofs" should not exist!

Non-uniform: witness hiding more formally

Let us formally give the verifier for UNSAT:

On input (y, A):

- Interpret A as a circuit
- Sample k tuples $(x_i, w_i) \sim \mathcal{D}$ and compute

$$p = (1/k) \sum_i \mathbb{1}[(x_i, A(x_i, \mathsf{NIWI}(x_i \vee y,\ w_i)) \in R_L]$$

- Accept iff p is sufficiently large

Non-uniform: technical conditions

Because verifier is randomized, really a Merlin-Arthur proof system

Lots of technical issues related to asymptotics

- Allow verifier slightly super-polynomial runtime, witness length
- Must assume verifier fails on all but finitely many input lengths
- Need NIWI, L search problem super-polynomially hard

Table of Contents

Background

Outline of results

Technical details

1. Witness hiding in 2 messages
2. Non-uniform witness hiding
3. Universal non-interactive proofs
4. Witness hiding vs witness encryption

Universal proofs: prerequisites

We have been talking about proofs of membership in NP languages

Now we need something slightly different: a formal proof system $\mathcal{S}$ for statements about Turing machines

For concreteness, can use Peano arithmetic

Universal proofs: construction of Π_U

Let D a TM with inputs (x, z). Define a statement:

$$S_x = \exists (z, D, \pi) \in \{0,1\}^{\ell} \ : \ D \text{ accepts } (x, z)$$
$$\wedge \ \pi \text{ is an } \mathcal{S}\text{-proof that } D \text{ is a sound NP verifier for } L$$

Let τ be an $\mathcal{S}$-proof that V_L is sound for L. The prover will output

NIWI
"S_x" witness: w, V_L, τ

Universal proofs: soundness

$$S_x = \exists (z, D, \pi) \in \{0,1\}^{\ell} \;:\; D \text{ accepts } (x, z)$$
$$\wedge \;\; \pi \text{ is an } \mathcal{S}\text{-proof that } D \text{ is a sound NP verifier for } L$$

If second clause is true, then D is sound for L

So if first clause is true, conclude $x \in L$

Universal proofs: witness hiding

Let $\Pi' = (P', V')$ be any NIWH scheme.

Let π be the $\mathcal{S}$-proof that V' is sound

NIWI		NIWI
"S_x"	$\cong$	"S_x"
witness: w, V_L, τ		witness: $P'(x, w), V', \pi$

So given an attacker against Π_U, we can build an attacker against Π' by switching to the right-hand proof. Thus Π_U is witness hiding

Universal proofs: inefficient and non-uniform Π'

Since P', V' are not used in the actual construction
 P' can be inefficient and both can be non-uniform

However, the proof of correctness π must prove soundness for
 a particular choice of advice

Since our non-uniform construction does not have this property
 it does not suffice to show the universal scheme works

Universal proofs: other properties

Did not use anything special about witness hiding in security proof

In fact the same proof should go through
for any falsifiable security property

We claim this scheme is the best possible non-interactive proof

Table of Contents

Background

Outline of results

Technical details

1. Witness hiding in 2 messages
2. Non-uniform witness hiding
3. Universal non-interactive proofs
4. Witness hiding vs witness encryption

Witness hiding vs witness encryption: definitions

Witness encryption: encryption where x serves as public key, and L-witness w serves as private key

Formally two properties:

Correct: $\forall m \in \{0,1\}$, $(x, w) \in R_L$:

$$\mathrm{dec}(x, w, \mathrm{enc}(x, m)) = m$$

Soundness secure: $\forall A$ ppt, $x \notin L$, $m \sim \{0,1\}$:

$$\Pr[A(\mathrm{enc}(x, m)) = m] = \frac{1}{2} + \mathrm{negl}$$

Only known from strong tools (e.g. iO)

Witness hiding vs witness encryption: definitions

Weaker average case notion of correctness

For infinitely many security parameters and some polynomial p,

$$\Pr[\mathrm{dec}(x, w, \mathrm{enc}(x, m)) = m] = 1/p$$

and otherwise dec outputs $\perp$

Probability taken over choice of $(x, w) \sim T$
and internal randomness of both algorithms

Witness hiding vs witness encryption: construction

Fix $T \in \mathsf{NP} \cap \mathsf{coNP}$ with $\mathcal{E}$ a distribution over $(y, z) \in R_T$

Prover:

sample $(y, z) \sim \mathcal{E}$

compute NIWI π of $x \vee (y \notin T)$ using witness w

output y, z, π

Verifier: check π is valid and $(y, z) \in R_T$

Witness hiding vs witness encryption: soundness

$$x \vee (y \notin T)$$

$(y, z) \in R_T$ implies that $y \in T$

Conclude x must be true

Witness hiding vs witness encryption: witness hiding

Let A an adversary against witness hiding

Propose a witness encryption scheme. Instead of directly encrypting a message, we encrypt a randomly chosen value w.

$$
\begin{array}{ll}
\text{enc}(y, m): & \text{sample } (x, w) \sim \mathcal{D} \\
& \text{compute NIWI } \pi \text{ of } x \vee (y \notin T) \text{ using witness } w \\
& \text{output } (x, \pi) \\
\text{dec}(y, z, (x, \pi)): & \text{run } A(x, y, z, \pi) \text{ to get } w' \\
& \text{if } (x, w') \notin R_L \text{ output } \perp \text{, otherwise output } w'
\end{array}
$$

As L has unique witnesses know $w' = w$ when A succeeds

To encrypt a choosen bit m output $r, \langle w, r \rangle \oplus m$

Conclusion

Consider these four schemes as evidence that NIWH should exist

At the very least they are strong barriers to proving otherwise!

Reference I

Boaz Barak, Yehuda Lindell, and Salil P. Vadhan.
Lower bounds for non-black-box zero knowledge.
In *44th FOCS*, pp. 384–393, October 2003.

Boaz Barak, Shien Jin Ong, and Salil P. Vadhan.
Derandomization in cryptography.
In *CRYPTO 2003*, pp. 299–315, August 2003.

Ran Canetti, Yilei Chen, Justin Holmgren, Alex Lombardi, Guy N. Rothblum, Ron D. Rothblum, and Daniel Wichs.
Fiat-Shamir: from practice to theory.
In *51st ACM STOC*, pp. 1082–1090, June 2019.

Cynthia Dwork and Moni Naor.
Zaps and their applications.
In *41st FOCS*, pp. 283–293, November 2000.

Uriel Feige, Dror Lapidot, and Adi Shamir.
Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract).
In *31st FOCS*, pp. 308–317, October 1990.

Uriel Feige and Adi Shamir.
Witness indistinguishable and witness hiding protocols.
In *22nd ACM STOC*, pp. 416–426, May 1990.

Shafi Goldwasser, Silvio Micali, and Charles Rackoff.
The knowledge complexity of interactive proof-systems (extended abstract).
In *17th ACM STOC*, pp. 291–304, May 1985.

Reference II

Oded Goldreich and Yair Oren.
Definitions and properties of zero-knowledge proof systems.
Journal of Cryptology, 7(1):1–32, December 1994.

Abhishek Jain, Yael Tauman Kalai, Dakshita Khurana, and Ron Rothblum.
Distinguisher-dependent simulation in two rounds and its applications.
In *CRYPTO 2017, Part II*, pp. 158–189, August 2017.

Rafael Pass.
On deniability in the common reference string and random oracle model.
In *CRYPTO 2003*, pp. 316–337, August 2003.

Chris Peikert and Sina Shiehian.
Noninteractive zero knowledge for NP from (plain) learning with errors.
In *CRYPTO 2019, Part I*, pp. 89–114, August 2019.