Black-box use of One-way Functions is Useless for Optimal Fair Coin-Tossing

Hemanta K. Maji   Mingyuan Wang

PURDUE UNIVERSITY

August, 2020 (CRYPTO–2020)

Two-party Fair Coin-tossing Protocol

- An input-less r-message interactive protocol where parties always agree on the output $b \in \{0, 1\}$ at protocol culmination
- Fairness requires that even if one party aborts during the execution of the protocol, the other party should still output a bit.
  - Every party maintains a defense coin, which is their output if the other party aborts.
- Insecurity is defined as how much an adversary can alter the expected output of the other party (compared to the honest execution).

Position Our Contribution among Prior Works

- In the Information-theoretic setting:
  - Any protocol is constant-insecure.
- Assuming the existence of One-Way Functions:
  - We have an explicit $\Theta\left(1/\sqrt{r}\right)$-insecure protocol. Blum (COMPCON–82), Broder-Dolev (FOCS-84), Awerbuch-Blum-Chor-Goldwasser-Micali (1985), Cleve (STOC-86)
- Assuming the existence of Oblivious Transfer:
  - We have an explicit $\Theta\left(1/r\right)$-insecure protocol. Gordon-Hazay-Katz-Lindell (STOC-08), Moran-Naor-Segev (TCC-09)
  - $\Theta\left(1/r\right)$-insecurity is unavoidable as Cleve (STOC-86) proves that any r-message protocol is $\Omega\left(1/r\right)$-insecure.

Can we construct optimal fair coin-tossing protocol based on one-way functions alone?

Our Contribution

Any black-box construction of fair coin-tossing protocol from one-way functions is $\Omega\left(1/\sqrt{r}\right)$-insecure.

- The protocols from the 1980s are optimal!
- We prove this result by extending the potential-based argument introduced by recent works (Khorasgani-Maji-Mukherjee (TCC-19), Khorasgani-Maji-Wang (2020)) to our setting.

Formulating the Problem

- We consider a fair coin-tossing protocol, where parties exchange a total of r messages.
- The expected output is X, refer to as bias-X protocol.
- Alice maintains a defense coin $\in \{0,1\}$, which is her output if Bob aborts.
    - She might update her defense when she prepares a new message, i.e., setting up a new defense.
- Bob set up his defense coin $\in \{0,1\}$ analogously.

For simplicity, we shall only consider fail-stop adversaries. That is, the adversary follows the protocol honestly, but may abort prematurely.

1. This weaker adversary is already powerful enough to do the most devastating attack.
2. Private-key cryptographic primitives (e.g., commitment schemes) suffice to ensure honest behavior.

Cleve's Negative Result

Cleve (STOC-86) showed that for any r-message protocol, there exists a computationally efficient (fail-stop) adversary that alter the expected output by $\Omega(1/r)$

- Hence, every r-message fair coin-tossing protocol is $\Omega(1/r)$-insecure, regardless of what hardness assumption one assumes

Definition

An r-message fair coin-tossing protocol is called an optimal fair coin-tossing protocol if it is $\mathcal{O}(1/r)$-insecure.

Known Positive Results

Assume the existence of one-way functions,

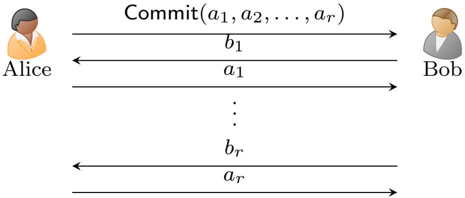

Alice samples private randomness $a_i \leftarrow \{0,1\}$.
Bob samples private randomness $b_i \leftarrow \{0,1\}$.
The output is the majority of

$$\begin{matrix} a_1 \oplus b_1 \\ a_2 \oplus b_2 \\ \vdots \\ a_r \oplus b_r \end{matrix}$$

- Majority protocol (Awerbuch-Blum-Chor-Goldwasser-Micali (1985), Cleve (STOC-86)), is $\Theta\left(1/\sqrt{r}\right)$-insecure.

. . . .

Assume the existence of oblivious transfer,

- Moran-Naor-Segev (TCC-09) constructed the optimal fair coin-tossing protocol, i.e., the MNS protocol is $\Theta\left(1/r\right)$-insecure.

Motivation

Summary of the state-of-the-art constructions:

	MNS Protocol		Majority Protocol
Assumption	Oblivious Transfer	stronger than	One-way function
Insecurity	$\Theta(1/r)$ (optimal)	more secure than	$\Theta(1/\sqrt{r})$

- In theoretical cryptography, a guiding principle is to build primitives using the minimal/weakest hardness of computation assumptions.
  - And if such constructions do not exist, what are the inherent hurdles?

Question

Can we construct optimal fair coin-tossing protocols from one-way functions
or
can we prove that it is inherently impossible?

- Unfortunately, one cannot prove the negative result unconditionally.
- One prominent technique of studying such questions is through the lens of black-box constructions (Impagliazzo-Rudich (STOC-89), Reingold-Trevisan-Vadhan (TCC-04)).

Black-box Constructions & Separations

A construction is (fully) black-box if the construction and the security reduction treat the primitive and the adversary in a black-box manner (Impagliazzo-Rudich (STOC-89), Reingold-Trevisan-Vadhan (TCC-04)).

Impagliazzo's World (Impagliazzo (CCC-95))

Minicrypt	Cryptomania
Pseudorandom Generator Håstad-Impagliazzo-Levin-Luby (1999) Commitment Scheme Naor (1991), Haitner-Reingold (STOC-07) Signature Scheme Rompel (STOC-90) ⋮	Key Agreement & Public-key Encryption Impagliazzo-Rudich (STOC-89) Oblivious Transfer Gertner-Kannan-Malkin-Reingold-Viswanathan (FOCS-00) ⋮

Whether optimal fair coin-tossing belongs to Minicrypt or Cryptomania remains one of the major open problems.

Our Results

Theorem (Informal)

Every black-box construction of an r-message bias-X fair coin-tossing protocol from one-way functions is $\Omega(X(1-X)/\sqrt{r})$-insecure.

Corollary (Implication 1)

Black-box use of one-way functions cannot yield optimal fair coin-tossing protocols.

Corollary (Implication 2)

Majority protocol is qualitatively the most secure protocol that one can build using one-way functions in a black-box manner.

Coin-tossing in the Random Oracle Model

Following the paradigm proposed by Impagliazzo-Rudich (STOC-89), we consider the coin-tossing protocols in the random oracle model.

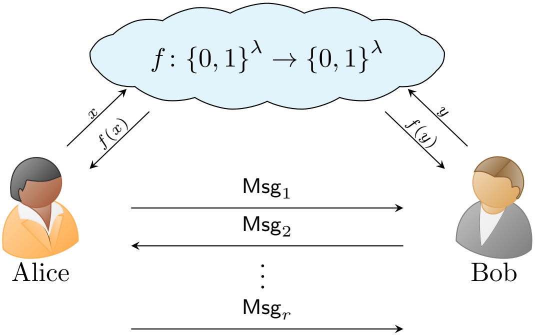

- Alice and Bob are computationally unbounded.
- In an honest execution, Alice and Bob ask $\mathsf{poly}(\lambda)$ queries. An adversary may ask additional queries to the random oracle.
- Intuitively, this models the usefulness of the black-box access to an "idealized" one-way function.
- **Objective.** We shall prove that there exists a fail-stop strategy that asks (at most) $\mathsf{poly}(\lambda)$ additional queries and alters the expected output by $\Omega(1/\sqrt{r})$.

Prior Works on Coin-tossing in the Random Oracle Model

- Dachman–Soled-Lindell-Mahmoody-Malkin (TCC-11) proved that if the message complexity $r = o\left(\frac{\lambda}{\log \lambda}\right)$, a fail-stop adversary can alter the expected output by $\Omega(1/\sqrt{r})$ by asking $2^{o(\lambda)}$ additional queries.
- Dachman–Soled-Mahmoody-Malkin (TCC-14) proved that if the protocol satisfies a special property called "function-oblivious", a fail-stop adversary can alter the expected output by $\omega(1/r)$ by asking $\mathsf{poly}(\lambda)$ additional queries.
  - Intuitively, "function-oblivious" requires that the output depends solely on the private randomness of each party; but is independent of the instantiation of the random oracle.
  - All the known protocols (e.g., majority protocols) are "function-oblivious".

. . . .

In comparison, our results resolve this problem in the full generality.

- We impose no restrictions on the message complexity or the type of protocols.
- The adversary asks polynomially many additional queries.
- The insecurity $\Omega(1/\sqrt{r})$ matches the positive result (Majority Protocol).
- Our results work for bias-X protocol with arbitrary $X \in (0, 1)$. X may depend on the security parameter.

Additional Relevant Work

In a sequence of work (Haitner-Nissim-Omri-Shaltiel-Silbak (FOCS-18), Haitner-Makriyannis-Omri (TCC-18)), Haitner, Makriyannis, and Omri proved that

- There exists a universal constant c, such that for any constant r, the existence of r-message fair coin-tossing protocol with insecurity $< c/\sqrt{r}$ implies the existence of (infinitely-often) key agreement protocols.

Comparison to this work

This work is incomparable to our results as it proves a stronger consequence but for restricted class of protocols.

Our Technical Proof

- Recall that we have a r-message bias-X fair coin-tossing protocol in the random oracle model.
- Our objective is to find a fail-stop adversary that asks (at most) $\mathsf{poly}(\lambda)$ additional queries and alters the expected output by $\Omega(1/\sqrt{r})$.

Correlation in the Random Oracle Model

Conditioned on the public transcript, Alice and Bob private views are correlated due to common private queries to the random oracle.

We shall first make Alice and Bob private views independent!

Heavy Querier

Heavy querier (Impagliazzo-Rudich (STOC-89), and Barak-Mahmoody (CRYPTO-09)) is a standard technique for removing correlations between Alice and Bob private view.

1. Public algorithm that takes the partial transcript as input and outputs a number of query/answer pairs
2. Guarantees that conditioned on partial transcript and Heavy querier's message, Alice and Bob private view are close to being independent
3. Asks polynomially many additional queries

Augmented Protocol

Immediately after every protocol message, the heavy querier is invoked and its message is attached.

- Note that this does not change the message complexity r.

Our Perspective

For every partial transcript v, define

- p_v is the probability that v happens
- x_v is the expected output conditioned on v
- a_v, b_v are the expectations of Alice's and Bob's defense coin conditioned on v

....

We are interested in finding a stopping time τ and the following score.

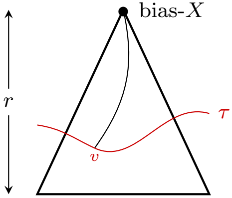

$$\mathsf{Score}(\tau) := \mathop{\mathrm{E}}_{v \leftarrow \tau}\left[|a_v - x_v| \; + \; |b_v - x_v| \right].$$

- $|a_v - x_v|$ is the change in Alice's expected output if Bob aborts at v.
- Analogously, $|b_v - x_v|$ is the change in Bob's expected output if Alice aborts.
- This score reflects the change in expected output when parties abort at τ
- We shall prove that $\max_\tau$ $\mathsf{Score}(\tau)$ is large.

An Inductive Approach

Following the recent work of Khorasgani-Maji-Mukherjee (TCC-19) and Khorasgani-Maji-Wang (2020), we use an inductive approach to prove that there exists a universal constant c such that

$$\max_{\tau} \mathsf{Score}(\tau) \geqslant c \cdot X(1-X)/\sqrt{r}.$$

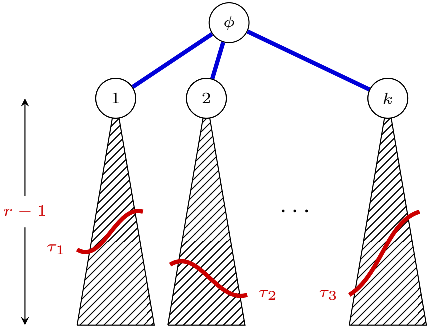

By our inductive hypothesis, $\max_\tau \mathsf{Score}(\tau)$ is higher than

$$p_1 \cdot \max\left(\overbrace{|a_1 - x_1| + |b_1 - x_1|}^{\text{Pick node 1 as stopping time}}, \overbrace{c \cdot x_1(1-x_1)/\sqrt{r-1}}^{\text{Inductive hypothesis}}\right)$$

$$+ p_2 \cdot \max\left(|a_2 - x_2| + |b_2 - x_2|, c \cdot x_2(1-x_2)/\sqrt{r-1}\right)$$

$$+ \cdots$$

$$+ p_k \cdot \max\left(|a_k - x_k| + |b_k - x_k|, c \cdot x_k(1-x_k)/\sqrt{r-1}\right)$$

$$= \mathop{\mathrm{E}}_{I}\left[\max\left(|a_I - x_I| + |b_I - x_I|, c \cdot x_I(1-x_I)/\sqrt{r-1}\right)\right]$$

The Potential Function

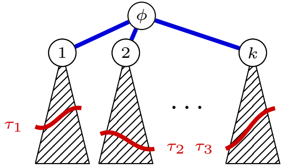

Now we need to prove

$$\mathop{\mathrm{E}}_{I}\left[\max\left(|a_I - x_I| + |b_I - x_I|\ ,\ c \cdot x_I\,(1 - x_I)\,/\sqrt{r-1}\right)\right] \geqslant c \cdot X(1-X)/\sqrt{r}$$

Khorasgani-Maji-Wang (2020) identified the following potential function

$$\boxed{\Phi(x,a,b) := x(1-x) + (x-a)^2 + (x-b)^2}$$

- $x(1-x)$ is the quality of the attack attributed to the bias of the protocol
- $(x-a)^2$ punishes Alice if her defense is too far away from the expected output
- $(x-b)^2$ punishes Bob if his defense is too far away from the expected output

They proved that

$$\max\left(|a_I - x_I| + |b_I - x_I|\ ,\ \frac{c}{\sqrt{r-1}} \cdot x_I\,(1-x_I)\right) \geqslant \frac{c}{\sqrt{r}} \cdot \Phi(x_I, a_I, b_I).$$

There could exist other potential functions. It happens to be the case that this one serves our purposes!

The Potential Function

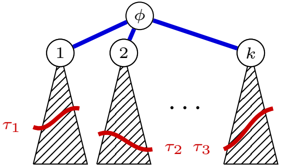

$$\Phi(x, a, b) := x(1-x) + (x-a)^2 + (x-b)^2$$

Our task reduces to proving

$$\mathop{\mathrm{E}}_{I}\left[\frac{c}{\sqrt{r}} \cdot \Phi(x_I, a_I, b_I)\right] \geqslant c \cdot X\left(1-X\right)/\sqrt{r}.$$

It suffices to prove that

$$\mathop{\mathrm{E}}_{I}\left[\Phi(x_I, a_I, b_I)\right] \geqslant X(1-X).$$

Completing the Proof

$\Phi(x,a,b) := x(1-x) + (x-a)^2 + (x-b)^2 = x + (x-a-b)^2 - 2ab$

$$
\begin{aligned}
\mathop{\mathrm{E}}_{I}\left[\Phi(x_I,a_I,b_I)\right] &= \mathop{\mathrm{E}}_{I}\left[x_I + (x_I-a_I-b_I)^2 - 2a_I b_I\right] \\
&= \mathop{\mathrm{E}}_{I}[x_I] + \mathop{\mathrm{E}}_{I}\left[(x_I-a_I-b_I)^2\right] - 2\cdot\mathop{\mathrm{E}}_{I}[a_I\cdot b_I] \\
&\geqslant \mathop{\mathrm{E}}_{I}[x_I] + \left(\mathop{\mathrm{E}}_{I}[x_I] - \mathop{\mathrm{E}}_{I}[a_I] - \mathop{\mathrm{E}}_{I}[b_I]\right)^2 - 2\cdot\mathop{\mathrm{E}}_{I}[a_I\cdot b_I]
\end{aligned}
$$

$\mathop{\mathrm{E}}_{I}[a_I b_I] = \mathop{\mathrm{E}}_{I}[a_I]\cdot\mathop{\mathrm{E}}_{I}[b_I]$ because Alice and Bob view are independent.

$$
\begin{aligned}
&= \mathop{\mathrm{E}}_{I}[x_I] + \left(\mathop{\mathrm{E}}_{I}[x_I] - \mathop{\mathrm{E}}_{I}[a_I] - \mathop{\mathrm{E}}_{I}[b_I]\right)^2 - 2\cdot\mathop{\mathrm{E}}_{I}[a_I]\cdot\mathop{\mathrm{E}}_{I}[b_I] \\
&= \Phi\left(\mathop{\mathrm{E}}_{I}[x_I], \mathop{\mathrm{E}}_{I}[a_I], \mathop{\mathrm{E}}_{I}[b_I]\right)
\end{aligned}
$$

Although $\Phi(x,a,b)$ is <u>not</u> a tri-variate convex function, we have identified a global invariant in the augmented protocols that ensures Jensen's inequality holds for this scenario.

Completing the Proof

The proof follows from

$$\Phi\left(\mathop{\mathrm{E}}_{I}[x_I], \mathop{\mathrm{E}}_{I}[a_I], \mathop{\mathrm{E}}_{I}[b_I]\right) = \Phi\left(X, \mathop{\mathrm{E}}_{I}[a_I], \mathop{\mathrm{E}}_{I}[b_I]\right)$$
$$= X(1-X) + \left(X - \mathop{\mathrm{E}}_{I}[a_I]\right)^2 + \left(X - \mathop{\mathrm{E}}_{I}[b_I]\right)^2 \geqslant X(1-X)$$

Summary of the Proof

1. We consider an r-message coin-tossing protocol in the random oracle model.
2. We use heavy querier algorithm to kill the correlations between Alice and Bob private view. This step asks $\mathsf{poly}(\lambda)$ additional queries.
3. We use the an inductive approach with a carefully crafted potential function to identify a stopping time τ, such that

$$\mathsf{Score}(\tau) := \mathop{\mathrm{E}}_{v \leftarrow \tau} \left[|a_v - x_v| + |b_v - x_v| \right] = \Omega\left(\frac{X(1-X)}{\sqrt{r}} \right)$$

4. This stopping time τ shall be translated into an attack that alter the expected output by $\Omega\left(\frac{X(1-X)}{\sqrt{r}}\right)$. When $X = 1/2$, this is $\Omega\left(\frac{1}{\sqrt{r}}\right)$.

Ongoing Work

We prove that optimal fair coin-tossing is also black-box separated from public-key encryption schemes.

Thanks!