On the Impossibility of Information-Theoretic Composable Coin Toss Extension

Gregor Seiler Ueli Maurer
Department of Computer Science
ETH Zürich, Switzerland
Email: {gseiler,maurer}@inf.ethz.ch

***Abstract*—Shared randomness is an important resource in cryptography. It is well-known that in the information-theoretic setting there is no protocol that allows two parties who do not trust each other to obtain a uniformly distributed shared bit string solely by exchanging messages such that a dishonest party can not influence the result. On the other hand, in the situation where the two parties already share a random bit string and want to use it in order to construct a longer random bit string, it is only known to be impossible when the protocols are restricted in the number of messages to be exchanged. In this paper we prove that it is also impossible when arbitrarily many messages are allowed.**

I. Introduction

The two-party task of flipping coins asks for protocols that enable two parties to generate a shared uniformly distributed bit string such that even if one party is dishonest and deviates arbitrarily from the protocol, it can not influence the distribution of the resulting bit string. This is important for instance when playing games over the internet. The first such protocol was given by Manuel Blum [1]. It is secure in a computational stand-alone security model. We focus on the information-theoretic setting and work in the constructive cryptography framework where security is maintained under arbitrary composition [2], [3]. In this security framework it is easily seen to be impossible to flip coins from scratch just by communication. Therefore, it is interesting to consider the scenario where the two parties already dispose of shared coins and want to obtain more. This was termed coin toss *extension* by Hofheinz et al. in [4]. See also [5] for a related but more general problem, albeit with weaker bounds. Hofheinz et al. studied the problem with respect to the computational and the information-theoretic flavour of both a stand-alone and the universal composability security notion and gave a nearly complete characterization of when coin toss extension is possible. It is in the case of information-theoretic composable security where their answer is not complete, for they show that information-theoretic composable coin toss extension is impossible with protocols that exchange less then a certain number of messages but not whether it is possible or not with protocols that proceed in arbitrarily many rounds. We close this gap by answering the question in the negative.

II. Preliminaries

A. Constructive Cryptography

In constructive cryptography every object is a system. In this paper we think about systems in their incarnation of random systems [6]. They have several interfaces each of which can be connected to an interface of another system so that the systems can exchange messages. The messages a particular system sends are controlled by conditional probability distributions given all earlier sent and received messages of that system. Now the basic building blocks of constructive cryptography are *resources*. These are systems that have one interface for every party. Since we remain in the two-party case, they have two interfaces in this paper. Resources provide the functionalities to the parties that are either assumed to be available or are aimed for in a protocol. For instance, a channel $\leftrightarrow$ is a resource which just forwards all messages from one party to the other. At the heart of constructive cryptography lies the notion of *constructing* resources from other resources in a precisely defined way with the help of *converters*. A converter is again a two-interface system. It can be plugged to the interface of a resource R belonging to a party A where it implements $A's$ part of a protocol by using the functionalities provided by R. It provides to A new functionalities at the free interface. In a sense, a converter can be seen to convert the interface of a resource with some functionalities to an interface with different functionalities. A protocol can be thought of as a pair of converters. Last there are *distinguishers* which have three interfaces. Two of the interfaces of a distinguisher D can be connected to the two interfaces of a resource R. D then outputs a bit at its third interface after communicating with R. So we obtain a random experiment $D(R)$ where the random variable Z that describes the output of D is defined. The goal for D is to distinguish R from another resource S by outputting 1 in the case of S and 0 in the case of R.

Definition 1. The *advantage* of a distinguisher D in distinguishing the resources R and S is given by

$$\Delta^D(R,S) = \Pr^{D(S)}(Z=1) - \Pr^{D(R)}(Z=1).$$

We write $R \approx_\varepsilon S$ if

$$\Delta^D(R,S) \leq \varepsilon$$

for all distinguishers D.

978-1-5090-1806-2/16/$31.00 ©2016 IEEE

In our case resources have two interfaces, which we consider as the left and right interface. We write terms like αR, which means that the converter α is connected to the left interface of the resource R. Similarly, we can write RS for the system we obtain when the right interface of R is connected to the left interface of S. Now we can formulate precisely what is understood by constructing a resource S from a resource R with the help of a protocol given by the converters α and β.

Definition 2. The protocol $\pi = (\alpha, \beta)$ *securely constructs* the ideal resource S from the available resource R with security parameter $\varepsilon > 0$ and both parties allowed to be dishonest if there are converters σ, τ such that

$$\begin{aligned} \alpha R \beta &\approx_\varepsilon S, \\ \alpha R \ &\approx_\varepsilon S\sigma, \\ R\beta &\approx_\varepsilon \tau S. \end{aligned}$$

The converters σ and τ are called simulators.

The first equation is the so-called correctness equation. It says that the protocol really constructs the desired resource. The second and third equation are the simulatability conditions. They model what a dishonest party is allowed to achieve. Namely, it is only allowed to achieve in the real world what it can also achieve in the ideal world with the help of a simulator. See [3, Theorem 1] for the fundamental theorem of constructive cryptography which says that the construction notion is maintained under composition.

B. Coin Tossing Resources

A 2-party coin tossing resource CT_n generates a uniformly random n-bit string, which can be received by both parties upon request. In this basic form it is trivially impossible to securely construct CT_n by any protocol from a weaker resource. This is because if there were such a protocol, then a distinguisher who does not participate at all in the protocol as party A, i.e., does not send any messages, would still need to receive the n-bit string for the other party B. The string would thus need to be generated by B alone. Similarly with A and B swapped so that the two strings of the two parties would differ with large probability contradicting the correctness condition. Therefore, we allow dishonest parties in the ideal world to obtain the n-bit string first and control when the other party receives it.

Definition 3. Let $n > m \geq 1$ be integers. A *coin tossing resource* CT_m with m coins is a two-party resource that generates a uniformly distributed m-bit string once and outputs it to both parties in response to every query. The resource CT'_n with n coins is like CT_n and, in addition, lets each party set a flag that has the following effect if set by party A, say. When party B queries the n-bit string, it is delivered to party A instead, who can then let the (unaltered) string be delivered to party B at any later point. The *filter* ϕ is a converter which shields away this flag so that $\phi CT'_n \phi = CT_n$.

We now formulate what we mean by a coin toss extension protocol. We assume that the parties do not only dispose of m coins, but also have a communication channel between them. So, we need to first introduce the concept of parallel composition.

Definition 4. The *parallel composition* $[R \mid S]$ of two resources R and S is the resource which at its left and right interfaces grants access to the corresponding interface of R and S.

Definition 5. An information-theoretic protocol for extending $m \geq 1$ to $n > m$ coins with security parameter ε is a protocol (α, β) such that

$$\alpha\,[CT_m \,|{\leftrightarrow}]\,\beta \approx_\varepsilon \phi CT'_n \phi \tag{1}$$

$$\alpha\,[CT_m \,|{\leftrightarrow}] \quad \approx_\varepsilon \phi CT'_n \sigma \tag{2}$$

$$[CT_m \,|{\leftrightarrow}]\,\beta \approx_\varepsilon \tau CT'_n \phi \tag{3}$$

with simulators σ, τ.

In this model both honest and dishonest parties can get the m-bit string at the beginning. This fits our intuition behind the problem, which is that the parties already have the m coins before they engage in a protocol to obtain more coins. One can also be interested in the situation where the parties receive the m coins after they have completed the protocol. Note that the results in [4] are proven in such a model. For honest parties this does not make a difference as they can just send in each round of the protocol each of the 2^m messages that they would have sent if they would have gotten the corresponding m-bit string in the beginning, and then decide on the set of messages in the end when they finally receive the coins. Dishonest parties on the other hand are potentially less powerful in this situation. This is reflected in that the simulators are not forced to commit to an m-bit string at the beginning, potentially allowing for a better simulation strategy. Therefore, an impossibility theorem for coin toss extension in such a model would be a stronger statement — it would establish impossibility even against weaker adversaries.

III. Impossibility Result

Theorem 1. *Let $n > m \geq 1$ be integers. There is no information-theoretic protocol for extending m to n coins with security parameter $\varepsilon < \frac{1}{8}$.*

Proof. Suppose that there are converters α, β and simulators σ, τ such that (1), (2) and (3) hold. Let D be the distinguisher that connects to n-bit coin tossing resources, fetches the random bit strings on both sides and outputs 1 precisely if both strings are equal. Connecting D to, for instance, CT_n yields the random experiment $D(CT_n)$, which defines the random variables $Y, Y' \in \{0,1\}^n$ and $Z \in \{0,1\}$ describing the bit strings received by D at the left and right interface of CT_n and the decision bit of D, respectively. Of course in this case Z is always equal to 1 by the definition of CT_n. We write

$$\Pr^{D(CT_n)}(Z = 1) = \Pr^{D(CT_n)}(Y = Y') = 1$$

for the probability that D outputs 1 in this experiment. Let χ be the system that connects on both sides to systems of

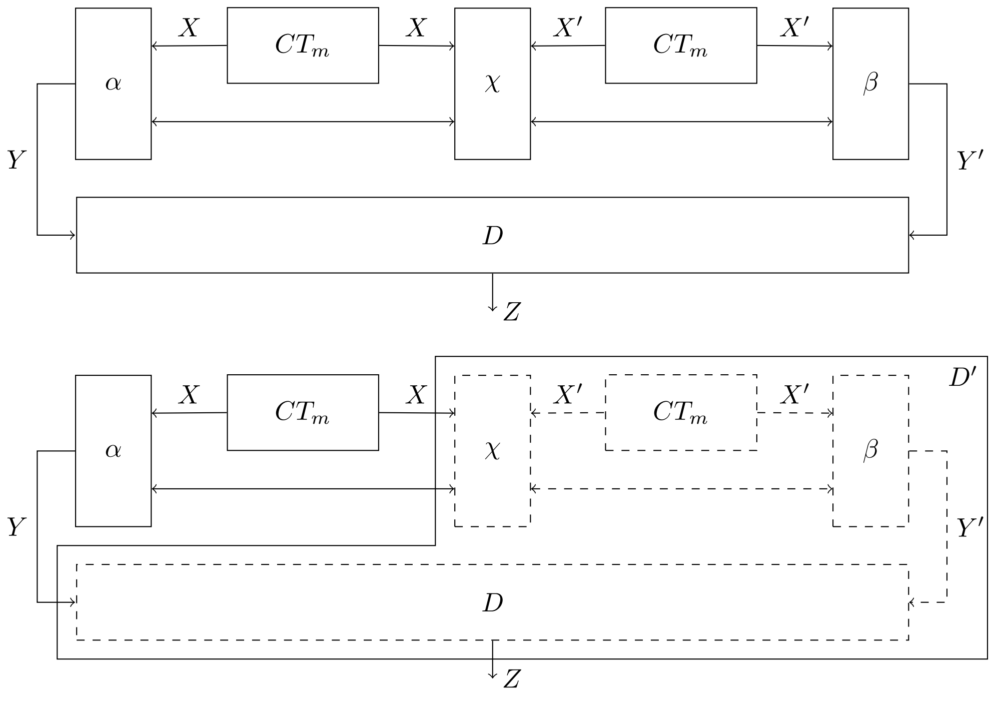

Figure 1: A graphical representation of the experiment E_2 in the proof of Theorem 1 together with the change in viewpoint when the right part of the system is absorbed into the distinguisher.

the type $[CT_m\,|{\leftrightarrow}]$, fetches the m-bit strings from both coin tossing resources and then forwards all messages between the channels. Now consider the experiments

$$D(\alpha[CT_m\,|{\leftrightarrow}]\beta), \tag{E_1}$$
$$D(\alpha[CT_m\,|{\leftrightarrow}]\chi[CT_m\,|{\leftrightarrow}]\beta), \tag{E_2}$$
$$D(\phi CT'_n\sigma\chi[CT_m\,|{\leftrightarrow}]\beta), \tag{E_3}$$
$$D(\phi CT'_n\sigma\chi\tau CT'_n\phi). \tag{E_4}$$

In each of the experiments there are the random variables $Y, Y' \in \{0,1\}^n$ and $Z \in \{0,1\}$ that describe the n-bit strings received by D at the left and right interface and the decision bit of D. Additionally, in the experiments E_2, E_3 and E_4 we have the random variables $X, X' \in \{0,1\}^m$ describing the m-bit strings received by χ at its left and right interface. For the first experiment E_1 we get from the correctness equation (1)

$$\begin{aligned}
&\Pr\nolimits^{E_1}(Z=1)\\
&= \Pr\nolimits^{D(CT_n)}(Z=1)\\
&\quad - (\Pr\nolimits^{D(CT_n)}(Z=1) - \Pr\nolimits^{D(\alpha[CT_m|{\leftrightarrow}]\beta)}(Z=1))\\
&= 1 - \Delta^D(\alpha[CT_m\,|{\leftrightarrow}]\beta, CT_n)\\
&\geq 1-\varepsilon.
\end{aligned} \tag{4}$$

In the experiment E_2 all messages sent by α and β and ultimately the n-bit strings Y and Y' are controlled by conditional probability distributions given earlier messages and the m-bit string from one of the CT_m, i.e., given X in the case of α and X' in the case β. Therefore, under the condition that $X = X'$, all probabilities in E_2 are equal to the corresponding probabilities in E_1. In particular

$$\Pr\nolimits^{E_2}(Z=1 \mid X=X') = \Pr\nolimits^{E_1}(Z=1) \geq 1-\varepsilon. \tag{5}$$

Next we will find with the help of the simulatability equation (2) that the conditional probabilities in E_3 given $X = X'$ are not much different from the corresponding probabilities in E_2. Concretely, we will show that

$$\begin{aligned}
&\Pr\nolimits^{E_2}(Z=1 \mid X=X')\\
&\leq \Pr\nolimits^{E_3}(Z=1 \mid X=X') + \varepsilon.
\end{aligned} \tag{6}$$

This equation together with Equation 5 reflects the intuition that the goal of the simulator σ in E_3 is to play the protocol in such a way as to influence the resulting n-bit string Y' so that it is equal to the string Y it has received from CT'_n. To go from E_2 to E_3 we absorb the right part $\chi[CT_m\,|{\leftrightarrow}]\beta$ of the systems in E_2 and E_3 into the distinguisher D in order to get a distinguisher for the left parts $\alpha[CT_m\,|{\leftrightarrow}]$ and $\phi CT'_n\sigma$ alone. See Figure 1 for a graphical representation of this change in viewpoint. Then the distinguisher knows the m-bit string of $\alpha[CT_m\,|{\leftrightarrow}]$ respectively $\phi CT'_n\sigma$, say $x \in \{0,1\}^m$, before $[CT_m\,|{\leftrightarrow}]\beta$ sends or receives its first message. Therefore, instead of emulating the unaltered system $[CT_m\,|{\leftrightarrow}]\beta$ internally, which outputs a uniform $x' \in \{0,1\}^m$, it can sample the system from the instances that output x. The fact that in both experiments E_2 and E_3 the random variable X', which describes the m-bit string generated by the CT_m in the right part, is uniformly distributed, implies that X, which describes the m-bit string from $\alpha[CT_m\,|{\leftrightarrow}]$

respectively $\phi CT'_n\sigma$, is independent of $X = X'$ in these experiments. Therefore, the new distinguisher simulates the experiments E_2 or E_3 conditioned on $X = X'$ when it is connected to $\alpha[CT_m\ |{\leftrightarrow}]$ or $\phi CT'_n\sigma$, respectively. We call this distinguisher D' and get from Equation 2

$$\begin{aligned}&\Pr^{E_2}(Z=1\mid X=X')\\&=\Pr^{D'(\alpha[CT_m|\leftrightarrow])}(Z=1)\\&\le\Pr^{D'(\phi CT'_n\sigma)}(Z=1)+\varepsilon\\&=\Pr^{E_3}(Z=1\mid X=X')+\varepsilon,\end{aligned}$$

which proves Equation 6.

Next we want to substitute $\tau CT'_n\phi$ for the right part $[CT_m\ |{\leftrightarrow}]\beta$ of the system in E_3 in order to go from E_3 to E_4. The method will basically be the same as before. Absorbing an appropriately sampled version of the left part $\phi CT'_n\sigma\chi$ into the distinguisher will allow us to apply the simulatability equation (3). By decomposing the event $Z = 1$ into the subevents where the m-bit string from CT_m attains a specific value $X' = x'$, we can write

$$\begin{aligned}&\Pr^{E_3}(Z=1\mid X=X')\\&=\sum_{x'\in\{0,1\}^m}\Pr^{E_3}(X'=x'\mid X=X')\\&\qquad\qquad\cdot\Pr^{E_3}(Z=1\mid X=X'=x'). \qquad (7)\end{aligned}$$

Here $X' = x'$ is not independent of $X = X'$ since X is not necessarily uniformly distributed in the experiment E_3. From the fact that, on the other hand, X' is indeed uniformly distributed, we get

$$\begin{aligned}\Pr^{E_3}(X'=x'|X=X')&=\frac{\Pr^{E_3}(X'=x')\Pr^{E_3}(X=x')}{\Pr^{E_3}(X=X')}\\=\frac{\Pr^{E_3}(X'=x')\Pr^{E_3}(X=x')}{2^{-m}}&=\Pr^{E_3}(X=x').\end{aligned}$$

Now let D'' be the distinguisher that connects to $[CT_m\ |{\leftrightarrow}]\beta$ and tries to simulate the experiment E_3 with $X = X'$ by emulating the systems D, χ and, after it knows the m-bit string x' from CT_m, the system $\phi CT'_n\sigma$ sampled from the instances that output the same string x', unless $\Pr^{E_3}(X = x') = 0$ in which case there is no such instance. In this case D'' outputs $Z = 0$ immediately. The probability

$$\begin{aligned}&\Pr^{D''([CT_m|\leftrightarrow]\beta)}(Z=1)\\&=\sum_{x'\in\{0,1\}^m}\Pr^{E_3}(X'=x')\mathbf{1}_{>0}(\Pr^{E_3}(X=x'))\\&\qquad\qquad\cdot\Pr^{E_3}(Z=1\mid X=X'=x') \qquad (8)\end{aligned}$$

is not precisely equal to (7) since instead of $\Pr^{E_3}(X = x')$ there is the factor $2^{-m}\mathbf{1}_{>0}(\Pr^{E_3}(X = x'))$ where $\mathbf{1}$ is the indicator function. Fortunately, this amounts to an error of at most ε. Intuitively the reason is that Equation (2) implies that the distribution of X is not far from the uniform distribution. Consider the distinguisher for $\phi CT'_n\sigma$ and $\alpha[CT_m\ |{\leftrightarrow}]$ that retrieves the m-bit string, say $x' \in \{0,1\}^m$, and then outputs 1 with probability

$$\mathbf{1}_{>0}(\Pr^{E_3}(X=x'))\Pr^{E_3}(Z=1\mid X=X'=x').$$

This distinguisher outputs 1 precisely with probabilities (7) or (8) when connected to $\phi CT'_n\sigma$ or $\alpha[CT_m\ |{\leftrightarrow}]$, respectively. Hence it follows from the simulatability equations (2) and (3) that

$$\begin{aligned}&\Pr^{E_3}(Z=1\mid X=X')\\&\le\Pr^{D''([CT_m|\leftrightarrow]\beta)}(Z=1)+\varepsilon\\&\le\Pr^{D''(\tau CT'_n\phi)}(Z=1)+2\varepsilon. \qquad (9)\end{aligned}$$

Together Equations (5), (6) and (9) read

$$\Pr^{D''(\tau CT'_n\phi)}(Z=1)\ge 1-4\varepsilon.$$

On the other hand, we find

$$\begin{aligned}&\Pr^{D''(\tau CT'_n\phi)}(Z=1)\\&=\sum_{y'\in\{0,1\}^n}\sum_{x'\in\{0,1\}^m}\Pr^{E_4}(Y'=y')\\&\qquad\qquad\cdot\Pr^{E_4}(X'=x'\mid Y'=y')\\&\qquad\qquad\cdot\mathbf{1}_{>0}(\Pr^{E_4}(X=x'))\\&\qquad\qquad\cdot\Pr^{E_4}(Y=y'\mid X=x')\\&\le\frac{1}{2^n}\sum_{y'}\sum_{x'}\Pr^{E_4}(Y=y'\mid X=x')\\&=\frac{1}{2^n}\sum_{x'}\sum_{y'}\Pr^{E_4}(Y=y'\mid X=x')\\&=\frac{2^m}{2^n}\le\frac{1}{2}.\end{aligned}$$

It follows that

$$1-4\varepsilon\le\Pr^{D''(\tau CT'_n\phi)}(Z=1)\le\frac{1}{2}.$$

So, $4\varepsilon \ge \frac{1}{2}$ in contradiction to $\varepsilon < \frac{1}{8}$. □

References

[1] M. Blum, "Coin flipping by telephone a protocol for solving impossible problems," *SIGACT News*, vol. 15, no. 1, pp. 23–27, 1983.

[2] U. Maurer and R. Renner, "Abstract cryptography," in *Innovations in Computer Science — ICS 2011*, B. Chazelle, Ed. Tsinghua University Press, 2011, pp. 1–21.

[3] U. Maurer, "Constructive cryptography — a new paradigm for security definitions and proofs," in *Theory of Security and Applications*, ser. Lecture Notes in Computer Science, S. Mödersheim and C. Palamidessi, Eds., vol. 6993. Springer Berlin Heidelberg, 2012, pp. 33–56.

[4] D. Hofheinz, J. Müller-Quade, and D. Unruh, "On the (im-)possibility of extending coin toss," in *Advances in Cryptology — EUROCRYPT 2006*, ser. Lecture Notes in Computer Science, S. Vaudenay, Ed., vol. 4004. Springer Berlin Heidelberg, 2006, pp. 504–521.

[5] G. Demay and U. Maurer, "Common randomness amplification: A constructive view," in *2012 IEEE Information Theory Workshop, Lausanne, Switzerland*, 2012, pp. 35–39.

[6] U. Maurer, "Indistinguishability of random systems," in *Advances in Cryptology — EUROCRYPT 2002*, ser. Lecture Notes in Computer Science, L. Knudsen, Ed., vol. 2332. Springer Berlin Heidelberg, 2002, pp. 110–132.