A Simple and Efficiently Verifiable Characterization of the Possibility of Information-Theoretic Key Agreement Secure Against Active Adversaries

Stefan Wolf

Computer Science Department
Swiss Federal Institute of Technology (ETH Zürich)
CH-8092 Zürich, Switzerland
E-mail address: wolf@inf.ethz.ch

Abstract

The model of information-theoretic secret-key agreement from joint randomness by public discussion was recently extended to the case where the insecure communication is not even authentic. It has been shown that the ability of generating a virtually-secret key is then directly linked to a certain "simulatability" condition formulated in terms of the involved random variables. More generally, this condition is important in the context of identification and authentication among parties sharing some correlated but not necessarily identical partially-secret keys. Unfortunately, the simulatability condition is a priori not very useful since it is not even clear whether it is verifiable in finite time. We introduce a new intuitive formalism, based on a mechanical model for representing the involved quantities, for dealing with discrete joint distributions of random variables and their manipulations by noisy channels, and show that this representation leads to a simple and efficient characterization of the possibility of secret-key agreement against active adversaries in many cases. The formalism is useful also for solving different problems related to discrete distributions and channels, e.g., to give criteria for the possibility and impossibility of secret-key agreement in the presence of *passive* opponents.

Keywords. Cryptography, unconditional security, active adversaries, identification, authentication, key agreement, secret-key rate.

1 Secret-Key Agreement Secure Against Passive and Active Adversaries

This paper is concerned with information-theoretic security in cryptography or, more precisely, key agreement unconditionally secure against active adversaries. Generalizing earlier models by Wyner [12] and Csiszár and Körner [2] based on communication over noisy channels, Maurer [5] and subsequently Ahlswede and Csiszár [1] have proposed the following interactive model of secret-key agreement by *authenticated* public discussion from common information. The parties Alice and Bob who want to establish a mutual secret key have access to realizations of random variables X and Y, respectively, whereas the adversary knows a random variable Z. Let P_{XYZ} be the joint distribution of the random variables. An example of a possible physical implementation is a satellite sending random bits at low signal power that are received by the parties with certain errors. Furthermore, the legitimate partners are connected by an insecure but authentic channel, i.e., a channel that can be passively overheard by Eve but over which no undetected active attacks by the opponent, such as modifying or inserting messages, are possible (see Figure 1).

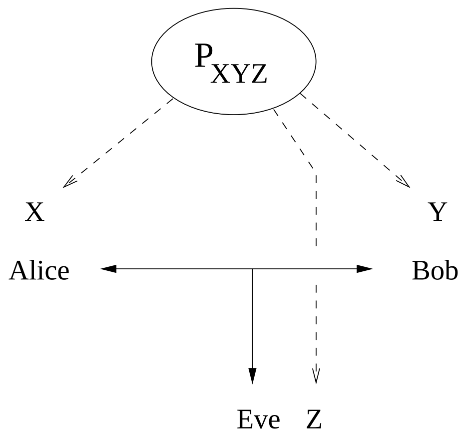

Figure 1: Secret-Key Agreement by Public Discussion from Common Information

In analogy to the models of Wyner and Csiszár-Körner, where the memoryless channels can be used many times, it is assumed here that the parties have access to many independent realizations of the corresponding random variables. The so-called *secret-key rate* in this model is the maximal rate at which Alice and Bob can generate a highly-secret key by communication over the insecure channel, measured with respect to the number of realizations of X and Y necessary for the generation of the key.

More precisely, the secret-key rate $S(X;Y||Z)$ of the joint distribution P_{XYZ} has been defined [5], [4], [10] as the maximal real number $R \geq 0$ with the property that for all $\varepsilon > 0$ and sufficiently large N, Alice and Bob can, by authenticated public communication, compute keys S_A and S_B from the blocks $X^N := [X_1, X_2, \ldots, X_N]$ and Y^N, respectively, such that S_A and S_B are both equal to a perfectly uniformly distributed key S with probability at least $1 - \varepsilon$, and such that

$$\log_2 |\mathcal{S}| \geq R - \varepsilon$$

holds (if $\mathcal{S}$ denotes the range of S), where the (Shannon) information about S provided by the communication C held over the public channel and by Eve's information Z^N must be at most ε, i.e.,

$$H(S|CZ^N) \geq \log_2 |\mathcal{S}| - \varepsilon \ .$$

Let us now consider secret-key agreement protocols that are supposed to be secure against *active* opponents. Clearly, one cannot expect that such a protocol is always successful if the adversary has full control over the public channel and can for instance block it completely, thus preventing any communication between Alice and Bob. Hence the best that can be achieved by such a protocol is that key agreement is successful when the adversary is passive, and that the parties realize failure due to an active attack and reject the outcome (see Figure 2).

To make things worse, we cannot even expect that a malicious active attack is always detected by both partners: Because Eve can always block the last (significant) message sent (the one that would make the second party accept), she can leave Alice and Bob in opposite acceptance states if this is her objective. However, nearly as strong robustness can indeed be defined and achieved. It can be required that (with high probability) either both Alice and Bob reject, or secret-key agreement is successful. Note that we cannot demand that both Alice and Bob accept in the latter case, but that they both compute the correct and secure key nonetheless.

According to this, the *secret-key rate* $S^*(X;Y||Z)$ *against active adversaries* (*robust secret-key rate* for short) was defined in [3], [11], [10] in the same way as $S(X;Y||Z)$, but where this time the public discussion channel is not even authentic (i.e., possibly under total control of the adversary). In addition to the conditions in the definition of $S(X;Y||Z)$, it is required that

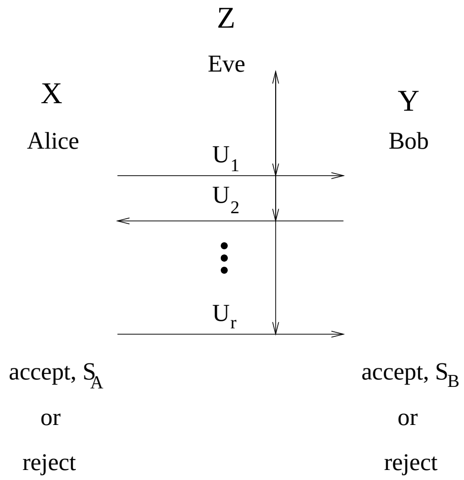

Figure 2: Unconditional Security Against Active Opponents

with probability at least $1 - \varepsilon$, either both parties reject the outcome of the protocol, or secret key agreement is successful.

2 The Simulatability Condition

Clearly, secret-key agreement secure against active adversaries as defined above can only be possible if Alice and Bob have some advantage over Eve in terms of the distribution P_{XYZ}. More precisely, it was shown that this advantage must be such that Eve cannot generate from Z a random variable $\overline{X}$ which Bob, knowing Y, is unable to distinguish from X (and vice versa). In [3], the following property of a distribution P_{XYZ} was defined.

Definition 1 [3] Let X, Y, and Z be random variables. Then X *is simulatable by* Z *with respect to* Y, denoted by

$$\mathrm{sim}_Y(Z \to X) ,$$

if there exists a conditional distribution $P_{\overline{X}|Z}$ such that $P_{\overline{X}Y} = P_{XY}$.

Another way of stating that $\mathrm{sim}_Y(Z \to X)$ holds is that there exists a random variable $\overline{X}$ such that $I(\overline{X}; XY|Z) = 0$, i.e., $XY \to Z \to \overline{X}$ is a Markov chain, with $P_{\overline{X}Y} = P_{XY}$.

Theorem 1 gives a complete characterization of the possibility of secret-key agreement against active adversaries, i.e., represents $S^*(X;Y||Z)$ in terms of P_{XYZ} and $S(X;Y||Z)$. The way of proving Theorem 1 that is sketched below (see [6]) is a simplified version of the proof of a similar result in [3].

Theorem 1 [3] *Let P_{XYZ} be a distribution. Then $S^*(X;Y||Z) = 0$ holds if either $\mathrm{sim}_Y(Z \to X)$ or $\mathrm{sim}_X(Z \to Y)$ holds. Otherwise, we have*

$$S^*(X;Y||Z) = S(X;Y||Z) \ .$$

Proof Idea. Clearly, $S^*(X;Y||Z) = 0$ holds when Eve can simulate one of the legitimate partners towards the other. On the other hand, if she is not able to simulate either X or Y, then key agreement at asymptotically the same rate as against an only passive wire-tapper is possible as follows. First, Alice and Bob generate an only short key by carrying out the protocol for the passive case, but authenticating each bit sent with a certain block of realizations of the random variable X or Y, respectively. (With a typical-sequences argument, one can show that Eve's success probability of an active attack can be made negligibly small.) Then a long secret key is generated by again using the protocol for the passive case, but this time the messages are authenticated by ε-almost-strongly-universal hashing [9], using the previously generated key. This way, the number of realizations of the random variables X and Y required for the authentication can be made asymptotically negligible as compared to the amount of randomness needed for the passive-adversary protocol. □

A result similar to the pessimistic implication of Theorem 1 (with the simulatability property as the important criterion) is even true in the scenario where the parties have access to only *one single* realization of the random variables [3], [10], [6]. More generally, simulatability is an important criterion for deciding whether an impersonation attack is possible in a scenario where parties are involved that share some correlated, but not perfectly equal, secret keys about which the opponent has some information.

3 A Calculus for Discrete Distributions and Channels

According to Theorem 1, the simulatability condition allows for separating the cases where secret-key agreement is possible and impossible in the presence of active adversaries. However, the characterization is a priori not practical because it depends on the existence of a particular channel (with certain properties) among the (uncountably-infinite) set of all discrete channels with given input and output alphabets. In the following, we hence consider the following questions:

- Is it, for a given distribution P_{XYZ}, possible to decide efficiently whether $\mathrm{sim}_Y(Z \to X)$ holds?
- Furthermore, if the answer is "yes," is it possible to efficiently find a channel $P_{\overline{X}|Z}$ such that $P_{\overline{X}Y} = P_{XY}$ holds?

We start by analyzing an example.

Example 1 Let the distribution P_{XYZ} of the random variables X, Y, and Z with ranges $\mathcal{X} = \{x_1, x_2\}$, $\mathcal{Y} = \{y_1, y_2\}$, and $\mathcal{Z} = \{z_1, z_2, z_3\}$ be as follows:

$$
\begin{array}{ll}
P_{XYZ}(x_1,y_1,z_1) = 6/100 \ , & P_{XYZ}(x_2,y_1,z_1) = 4/100 \ , \\
P_{XYZ}(x_1,y_1,z_2) = 9/100 \ , & P_{XYZ}(x_2,y_1,z_2) = 6/100 \ , \\
P_{XYZ}(x_1,y_1,z_3) = 15/100 \ , & P_{XYZ}(x_2,y_1,z_3) = 10/100 \ , \\
P_{XYZ}(x_1,y_2,z_1) = 36/100 \ , & P_{XYZ}(x_2,y_2,z_1) = 4/100 \ , \\
P_{XYZ}(x_1,y_2,z_2) = 9/100 \ , & P_{XYZ}(x_2,y_2,z_2) = 1/100 \ , \\
P_{XYZ}(x_1,y_2,z_3) = 0 \ , & P_{XYZ}(x_2,y_2,z_3) = 0 \ .
\end{array}
$$

In order to decide whether $\mathrm{sim}_Y(Z \to X)$ holds, we first consider the following representation of the (conditional) probabilities. We mark every symbol $x_i \in \mathcal{X}$ and every $z_j \in \mathcal{Z}$ with an empty or filled circle, where the size (or mass) of the circle corresponds to the probability $P_X(x_i)$ or $P_Z(z_j)$, and the position in the interval $[0,1]$ is given by the probability $P_{Y|X=x_i}(y_1)$ or $P_{Y|Z=z_j}(y_1)$, respectively (see Figure 3).

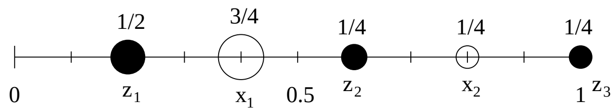

Figure 3: Representation of P_{XYZ}

Note first that not the entire information about P_{XYZ} is contained in this representation: only the distributions P_{XY} and P_{YZ}, but not P_{XYZ}, can be reconstructed from the quantities represented in the picture. We will see however that the fact whether or not X is simulatable by Z with respect to Y depends, not surprisingly, only on P_{XY} and P_{YZ}, as Theorem 2 shows. Second, not every representation corresponds to a distribution P_{XYZ}. This is only true if the total mass of each point set is 1, and if the marginal distribution P_Y is equal for both distributions P_{XY} and P_{YZ}. The last condition is equivalent to the fact that the sets of full and empty circles have the same center of gravity when interpreted as point masses.

Let now $Z^{(2)}$ with $\mathcal{Z}^{(2)} = \{z_1^{(2)}, z_2^{(2)}\}$ be generated by sending Z over the channel

$$
\begin{aligned}
P_{Z^{(2)}|Z}(z_1^{(2)}, z_1) &= 1 \, , \\
P_{Z^{(2)}|Z}(z_2^{(2)}, z_2) &= 1 \, , \\
P_{Z^{(2)}|Z}(z_2^{(2)}, z_3) &= 1 \, .
\end{aligned}
$$

For the new distribution $P_{XYZ^{(2)}}$, the above representation is as shown in Figure 4: Two masses have been united in their center of gravity.

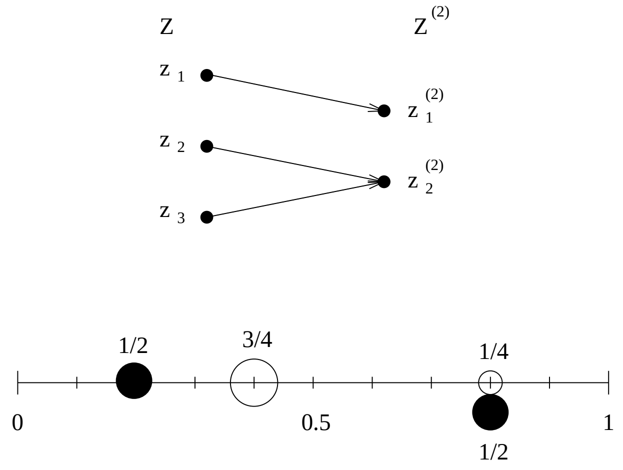

Figure 4: The Channel $P_{Z^{(2)}|Z}$, and $P_{XYZ^{(2)}}$

Let then $Z^{(2)}$ be sent over the additional channel $P_{Z^{(3)}|Z^{(2)}}$, where $\mathcal{Z}^{(3)} = \{z_1^{(3)}, z_2^{(3)}, z_3^{(3)}\}$, with

$$
\begin{aligned}
P_{Z^{(3)}|Z^{(2)}}(z_1^{(3)}, z_1^{(2)}) &= 1 \, , \\
P_{Z^{(3)}|Z^{(2)}}(z_2^{(3)}, z_2^{(2)}) &= 1/2 \, , \\
P_{Z^{(3)}|Z^{(2)}}(z_3^{(3)}, z_2^{(2)}) &= 1/2 \, .
\end{aligned}
$$

This corresponds to splitting one of the masses into two (equal) parts (see Figure 5).

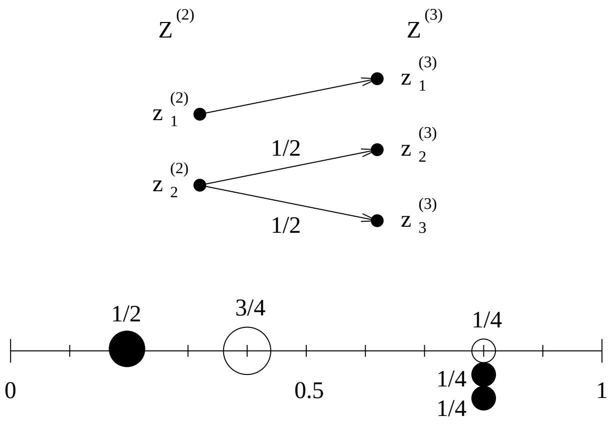

Figure 5: The Channel $P_{Z^{(3)}|Z^{(2)}}$, and $P_{XYZ^{(3)}}$

Finally, let $P_{\overline{X}|Z^{(3)}}$, with $\overline{\mathcal{X}} = \{\overline{x}_1, \overline{x}_2\}$, be given by

$$
\begin{aligned}
P_{\overline{X}|Z^{(3)}}(\overline{x}_1, z_1^{(3)}) &= 1 \ , \\
P_{\overline{X}|Z^{(3)}}(\overline{x}_1, z_2^{(3)}) &= 1 \ , \\
P_{\overline{X}|Z^{(3)}}(\overline{x}_2, z_3^{(3)}) &= 1 \ .
\end{aligned}
$$

The use of this channel again corresponds to uniting two masses in their center of gravity. The constellation of the masses with respect to X and $\overline{X}$ are now equal (see Figure 6), which means that $P_{\overline{X}Y} = P_{XY}$ holds. Hence $\mathrm{sim}_Y(Z \to X)$ is true, and the corresponding channel $P_{\overline{X}|Z}$ is the cascade of the three channels above:

$$
\begin{aligned}
P_{\overline{X}|Z}(\overline{x}_1, z_1) &= 1 \ , \\
P_{\overline{X}|Z}(\overline{x}_1, z_2) &= P_{\overline{X}|Z}(\overline{x}_2, z_2) = 1/2 \ , \\
P_{\overline{X}|Z}(\overline{x}_1, z_3) &= P_{\overline{X}|Z}(\overline{x}_2, z_3) = 1/2
\end{aligned}
$$

(see Figure 7).

We can now make this representation in the mechanical model more precise and exploit the direct connection between distributions and channels on one side and mass constellations as well as -operations on the other, in order to give a simple characterization of (non-)simulatability. The purpose of this physical model is to give more intuitive deductions and formulations of

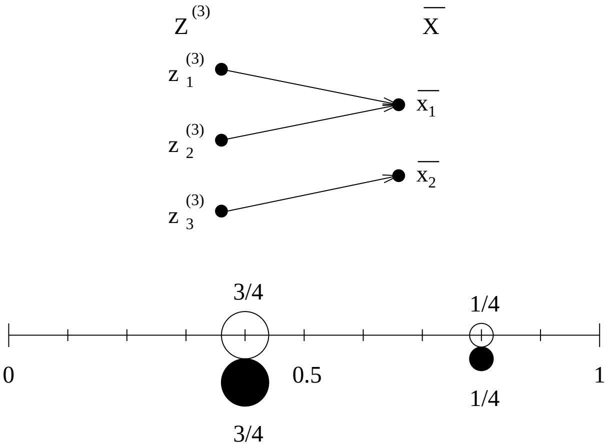

Figure 6: The Channel $P_{\overline{X}|Z^{(3)}}$, and $P_{XY\overline{X}}$

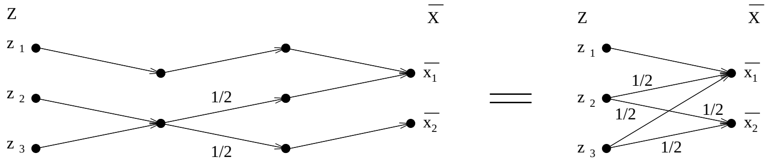

Figure 7: The Cascaded Channel $P_{\overline{X}|Z}$

results that could as well be stated and proved in terms of distributions and channels. Theorem 2 below makes a direct link between the two formalisms and justifies the point of view we take. In the following, particular emphasis lies on an intuitive presentation.

Definition 2 For an integer $N \geq 1$, an *N-dimensional (normed) mass constellation* $M := (m_i, a_i)_{i=1,\ldots,\ell}$ is a family of pairs with $m_i \in (0,1]$ and $a_i \in [0,1]^N$ for all i such that $\sum_i m_i = 1$. We additionally assume that the pairs are ordered with respect to the lexicographic ordering of the vectors a_i. The *center of gravity* (*center* for short) $c(M)$ of such a constellation is given by

$$c(M) := \sum_{i=1}^{\ell} m_i a_i \ .$$

Two constellations are *equicentered* if they have the same center of gravity. A constellation $M' = (m'_i, a'_i)_{i=1,\ldots,\ell'}$ is *derived* from $M = (m_i, a_i)_{i=1,\ldots,\ell}$ by *mass splitting* if $\ell' = \ell + 1$, and if there exist $0 < p < 1$, $1 \leq i_0 \leq \ell$, such that

$$(m'_i, a'_i) = \begin{cases} (m_i, a_i) & 1 \le i < i_0 \\ (pm_{i_0}, a_{i_0}) & i = i_0 \\ ((1-p)m_{i_0}, a_{i_0}) & i = i_0 + 1 \\ (m_{i-1}, a_{i-1}) & i_0 + 1 < i \le \ell + 1 \ . \end{cases}$$

Furthermore, M' is *derived* from M by *mass union* if $\ell' = \ell - 1$, and if there exist $i_1 < i_2$, $i_1 \le i_u \le i_2$, such that

$$(m'_i, a'_i) = \begin{cases} (m_i, a_i) & 1 \le i < i_1 \\ (m_{i+1}, a_{i+1}) & i_1 \le i < i_u \\ \left(m_{i_1} + m_{i_2}, \frac{m_{i_1} a_{i_1} + m_{i_2} a_{i_2}}{m_{i_1} + m_{i_2}}\right) & i = i_u \\ (m_i, a_i) & i_u < i < i_2 \\ (m_{i+1}, a_{i+1}) & i_2 \le i \le \ell - 1 \ . \end{cases}$$

We call mass splitting and union *basic mass operations*. A constellation M is called *stronger* than M', an event denoted by $M \rightsquigarrow M'$, if there exists a finite sequence of basic operations that transforms M into M'.

Let P_{UV} be the joint distribution of two random variables U and V with ranges $\mathcal{U}$ and $\mathcal{V} = \{v_1, \ldots, v_{N+1}\}$. Then the N-dimensional constellation $M_{U \leftarrow V}$ is defined by

$$M_{U \leftarrow V} = (P_U(u), (P_{V|U=u}(v_1), \ldots, P_{V|U=u}(v_N)))_{u \in \mathcal{U}} \ .$$

Note that the definition of $M_{U \leftarrow V}$ leads to a one-to-one correspondence between distributions P_{UV}, where $|\mathcal{V}| = N + 1$, and N-dimensional normed mass constellations $(m_i, a_i)_{i=1,\ldots,\ell}$ contained in the simplex characterized by $(a)_j \ge 0$ and $\sum_{j=1}^{N} (a)_j \le 1$.

It is clear that if $M \rightsquigarrow M'$, then the two constellations M and M' are equicentered. On the other hand, there exist equicentered constellations, none of which is stronger than the other (see Figure 8).

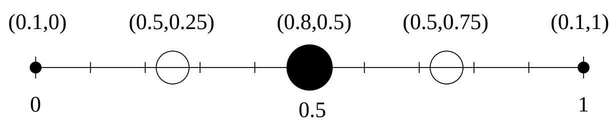

Figure 8: Incomparable Constellations: None is Stronger

Theorem 2 links simulatability and mass constellations. In this context, note first that for every distribution P_{XYZ}, $M_{X \leftarrow Y}$ and $M_{Z \leftarrow Y}$ are equicentered.

Theorem 2 *Let P_{XYZ} be the joint distribution of X, Y, and Z. Then X is simulatable by Z with respect to Y if and only if $M_{X\leftarrow Y}$ is stronger than $M_{Z\leftarrow Y}$:*

$$\mathrm{sim}_Y(Z \to X) \iff M_{Z\leftarrow Y} \rightsquigarrow M_{X\leftarrow Y} \ .$$

Proof. Let P_{U_1V} and P_{U_2V} be the joint distributions of random variables U_1 and V, and U_2 and V, respectively. Clearly, $M_{U_2\leftarrow V}$ can be obtained from $M_{U_1\leftarrow V}$ by a mass splitting or union operation if and only if there exists a "splitting channel" (as in Figure 4) or a "union channel" (see Figure 5) $P_{\overline{U}_2|U_1}$, respectively, such that

$$P_{\overline{U}_2V} = P_{U_1V} \cdot P_{\overline{U}_2|U_1} = P_{U_2V} \ .$$

The statement now follows from the facts that every discrete channel (with m output symbols) can be represented as a cascade of splitting and union channels, and that a cascade of channels is equivalent to the sequence of the corresponding mass operations. The first of these two facts can be shown as follows. First, all the letters of the input alphabet can be split, one after the other, to m symbols each (by $m-1$ splitting channels with certain probabilities for each symbol), and they can be united by union channels to the output symbols of the discrete channel. □

Remark. Note again that both conditions only depend on P_{XY} and P_{YZ}, but not otherwise on P_{XYZ}. Clearly, the condition given in Theorem 2 is a priori not more than a new formulation of simulatability, and is not obviously verifiable more efficiently. However, it leads to an efficiently-checkable criterion as Corollary 4 and Theorem 6 show.

As a preparation for the proof of Theorem 3, we describe a special mass operation, called *mass approach*, that can be composed by four basic operations (see Figure 9).

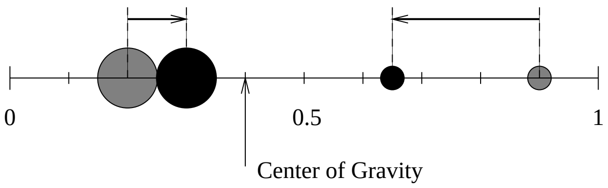

Figure 9: A Mass Approach

Lemma 1 *Let a constellation* $M = (m_i, a_i)_{i=1,\ldots,\ell}$ *be given, and let* $i \neq i'$, $1 \leq i, i' \leq \ell$. *We denote by*

$$c_{i,i'} := (m_i a_i + m_{i'} a_{i'})/(m_i + m_{i'})$$

the center of the i*th and* i'*th masses. Then for every* $\lambda \in [0,1]$, *there exist a sequence of four basic mass operations transforming* M *into the constellation that one obtains when the* i*th and* i'*th pairs are replaced by the pairs*

$$(m_i, a_i + \lambda(c_{i,i'} - a_i)) \quad \text{and} \quad (m_{i'}, a_{i'} + \lambda(c_{i,i'} - a_{i'}))$$

(which must be correctly put into the ordering).

Proof. The idea is that the masses m_i and $m_{i'}$ "exchange" a suitable mass $0 \leq m_e \leq \min\{m_i, m_{i'}\}$, i.e., that both masses are split into two parts, one of which is equal to m_e in both cases, and the union operation is applied twice to the remaining mass with the m_e-part of the other mass. Hence four basic operations are required. Because the choice $m_e = 0$ leaves m_i and $m_{i'}$ unchanged, whereas $m_e = \min\{m_i, m_{i'}\}$ corresponds to mass union, and since the result depends linearly on m_e, every position of m_i and $m_{i'}$ on their connecting line such that the masses are closer to each other, and such that the center of gravity remains unchanged, can be achieved this way. More explicitly, the mass m_e must be chosen as $\lambda \cdot \min\{m_i, m_{i'}\}$. □

We have now established the mechanical model and the necessary techniques for our characterizations of simulatability. In Corollary 4 we give a simple and efficiently verifiable, both necessary and sufficient condition for simulatability with respect to a *binary* random variable Y. Furthermore, the proof of Theorem 3 additionally shows that the corresponding channel $P_{\overline{X}|Z}$ can even be computed efficiently.

We first define what it means that a one-dimensional mass constellation is "more centered" than another. This relation leads to the characterization we are looking for. Note that this relation is not a linear ordering: When considering two random mass constellations, typically no one will be more centered than the other (see Figure 8).

Definition 3 For a one-dimensional mass constellation M and for $0 < t \leq 1$, we denote by $\ell_t(M)$ the leftmost masses of M of total amount d. (Typically, of one of the masses in M, only a part will be in $\ell_t(M)$.) A constellation M' is called *more centered* than M, denoted by

$$M' \prec M \ ,$$

if for all t,

$$c(\ell_t(M')) \geq c(\ell_t(M)) \ ,$$

where $c(S)$ stands for the center of gravity of a set S of masses.

Note first that this is a symmetric notion, i.e., that "left" and "$\geq$" could be replaced by "right" and "$\leq$" without changing the definition. Given two (finite) mass constellations, this quantity can be efficiently (i.e., in time linear in the total number of masses) checked. To see this, note that $M' = (m'_j, a'_j)_{j=1,\ldots,\ell'} \prec M$ is equivalent to the fact that for every $1 \leq k < \ell'$, the center of the set of masses $m'_1, \ldots, m'_k$ is not left of (i.e., smaller than) the center of $\ell_{m'_1+\cdots+m'_k}(M)$.

Theorem 3 *Let two equicentered one-dimensional mass constellations M and M' be given. Then M is stronger than M' if and only if M' is more centered than M:*

$$M \rightsquigarrow M' \iff M' \prec M \ .$$

Clearly, Corollary 4 follows immediately from Theorems 2 and 3.

Corollary 4 *Let P_{XYZ} be the joint distribution of random variables X, Y, and Z, where Y is binary. Then X is simulatable by Z with respect to Y if and only if $M_{X\leftarrow Y}$ is more centered than $M_{Z\leftarrow Y}$, i.e.,*

$$\mathrm{sim}_Y(Z \to X) \iff M_{X\leftarrow Y} \prec M_{Z\leftarrow Y} \ .$$

Proof of Theorem 3. We assume first that

$$M' = (m'_j, a'_j)_{j=1,\ldots,\ell'} \prec M = (m_i, a_i)_{i=1,\ldots,\ell}$$

holds. We show by induction that for every $0 \leq j_0 \leq \ell'$, there exists a sequence of basic mass operations that transforms M into a constellation $M_{j_0} = (\overline{m}_k, \overline{a}_k)_{k=1,\ldots,\overline{\ell}}$ such that for every $j \leq j_0$, there exists $k(j)$ (where $k(j) \neq k(j')$ if $j \neq j'$) with $\overline{m}_{k(j)} = m'_j$ and $\overline{a}_{k(j)} \leq a'_j$, and such that the center of the masses $\overline{m}_1, \ldots, \overline{m}_{j_0}$ of M_{j_0} is equal to $c(\ell_{\overline{m}_1+\cdots+\overline{m}_{j_0}}(M))$.

Clearly, this holds for $j_0 = 0$. We assume that the statement is true for $0 \leq j_0 < \ell'$ and show its validity also for $j_0 + 1$. Let $M_{j_0} = (\overline{m}_k, \overline{a}_k)_{k=1,\ldots,\overline{\ell}}$ be defined as above.

We transform M_{j_0} into M_{j_0+1} as follows. First, the leftmost among the masses $\overline{m}_{j_0+1}, \overline{m}_{j_0+2}, \ldots$, of total amount m'_{j_0+1}, are united into their center of gravity. Let $(\overline{m}_{j_0+1}, \overline{a}_{j_0+1}) = (m'_{j_0+1}, \overline{a}_{j_0+1})$ be the resulting mass. Then, because of $M' \prec M$ and by the induction assumption, the center of the masses $(\overline{m}_1, \overline{a}_1), \ldots, (\overline{m}_{j_0+1}, \overline{a}_{j_0+1})$ is not on the right-hand side of the center of gravity of $\ell_{m'_1+\ldots+m'_{j_0+1}}(M')$. Hence there exists a sequence of mass approaches, applied only to masses among $\overline{m}_1, \ldots, \overline{m}_{j_0+1}$, such that each of the the resulting masses (still of the same sizes) is on the left-hand side of (or at the same position as) the corresponding mass of M' (see Figure 10). Hence this new constellation satisfies the induction assumption for $j_0 + 1$, and this concludes the induction argument.

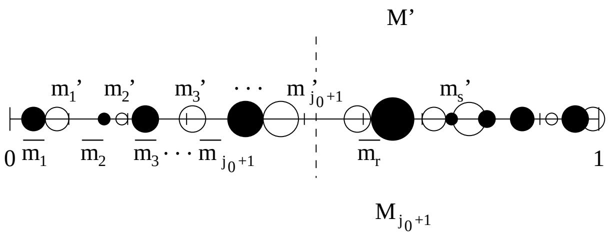

Figure 10: M' and M_{j_0+1}

Therefore M is stronger than some $\overline{M}$ satisfying the above property, with respect to M', for $j_0 = \ell'$. However, because $\overline{M}$ and M' are both equicentered to M, and because all masses of $\overline{M}$ lie, roughly speaking, on the left of (or at the same place as) the corresponding masses of M', we must have that $\overline{M} = M'$, hence $M \rightsquigarrow M'$.

We show the necessity of the condition. Assume for M and M' and for some t that

$$c(\ell_t(M')) < c(\ell_t(M)).$$

Then $M \not\rightsquigarrow M'$ holds because the basic mass operations, i.e., mass union (mass splitting leaves all the centers unchanged), can only move the center of the set $\ell_t(M)$ to the right (union of two masses, one in the set $\ell_t(M)$, and one in the complement) or leave it at the same place (union within $\ell_t(M)$ or within the complement). □

Note that the criterion for simulatability of Corollary 4 is not only very simple and efficiently verifiable, but that the proof of Theorem 3 also shows how a channel $P_{\overline{X}|Z}$ for simulating X with respect to Y can be constructed efficiently.

Let us now, after the complete analysis of the case of a binary random variable Y, consider the general case again. In Definition 4, we give a straight-forward, and also efficiently checkable, generalization of the notion that a constellation is more centered than another. This leads to a *necessary* criterion for simulatability (Theorem 6). However, although it appears to be sufficient as well in many cases, we give an example for which non-simulatability is not detected by the criterion.

Definition 4 Let M and M' be two N-dimensional mass constellations. Let furthermore a line L, passing through the origin, be given. We now consider the orthogonal projections of all the masses in the N-dimensional space onto L. This yields two one-dimensional equicentered mass constellations M_L and M'_L. We say that M' is *more centered* than M, $M' \prec M$, if $M'_L \prec M_L$ for every line L.

It is not difficult to see that also in N dimensions $M \rightsquigarrow M'$ can only hold if $M' \prec M$ holds. The reason is that $M_L \rightsquigarrow M'_L$ follows from $M \rightsquigarrow M'$: Projections of mass operations are mass operations again.

Theorem 5 *Let M and M' be N-dimensional equicentered mass constellations. If M is stronger than M', then M' must be more centered than M:*

$$M \rightsquigarrow M' \implies M' \prec M \ .$$

Corollary 6 *Let P_{XYZ} be the joint distribution of X, Y, and Z. If $M_{X\leftarrow Y}$ is* not *more centered than $M_{Z\leftarrow Y}$, then X is* not *simulatable by Z with respect to Y, i.e.,*

$$M_{X\leftarrow Y} \not\prec M_{Z\leftarrow Y} \implies (\,\mathrm{sim}_Y(Z \to X) \quad \textit{does } \text{not } \textit{hold}\,) \ .$$

Note that this condition is, despite the fact that the number of lines through the origin is infinite, efficiently verifiable since the number of points is finite. First, not every direction, i.e., every line, has to be checked separately. There are only at most $\binom{\ell+\ell'}{2}$ directions for which the mass constellations are different (with respect to the order of the masses), where ℓ and ℓ' are the numbers of masses in M and M', respectively. Equal orders means that, in the N-dimensional space, the same masses are "leftmost." Hence, all these directions can be treated simultaneously by looking at extremal directions. Furthermore, only the values t corresponding to a subset of the masses in M' have to be considered (as in the one-dimensional case).

Unfortunately, the given condition is not sufficient for simulatability (i.e., for a mass constellation being stronger than another) in the $N(\geq 2)$-dimensional case (although it appears to be a "good" condition failing to detect non-simulatability only in a small fraction of all cases). The following is a counterexample.

Example 2 Consider the following two-dimensional mass constellations M and M'.

$$\begin{aligned} M &= (0.2,(0,0)),\ (0.2,(0,0.5)),\ (0.2,(0.5,0)), \\ &\quad (0.2,(0.5,0.5)),\ (0.2,(0.25,0.25)), \\ \\ M' &= (0.2,(0.1,0)),\ (0.2,(0.1,0.5)),\ (0.2,(0.4,0)), \\ &\quad (0.2,(0.4,0.5)),\ (0.1,(0.15,0.25)),\ (0.1,(0.35,0.25)) \end{aligned}$$

(see Figure 11).

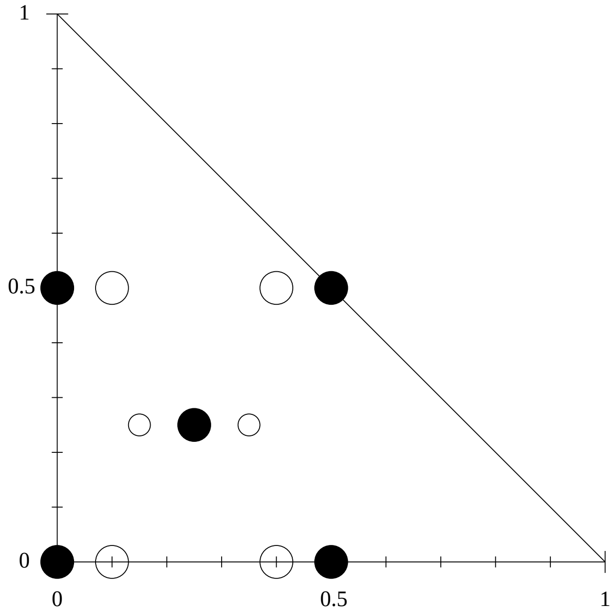

Figure 11: A 2-Dimensional Counterexample

It is not difficult to see that $M' \prec M$ holds. First, it clearly holds for the horizontal line and, because the (horizontal) distances between neighboring masses change in the same ratios, for all lines except the vertical line, for which the projected constellations are identical however.

On the other hand, M cannot be transformed into M' by basic operations. This is true because when considering the projection to the vertical line, it is clear that no union operation can be made except between masses with the same y-coordinate. However, with such operations only, M can clearly not be transformed to M' because of the masses with y-coordinate $1/2$. Hence $M \not\leadsto M'$ holds.

4 Concluding Remarks

We have analyzed the so-called simulatability condition which is of central importance in the context of unconditionally-secure identification and authentication between parties sharing some correlated information. For instance, this condition characterizes the possibility of secret-key agreement based on joint randomness in the presence of an active adversary. However, the criterion was not shown to be practical previously; it was not even clear whether it can be checked even in finite time.

We have introduced a new formalism for representing joint distributions of discrete random variables and their manipulations by noisy channels in a mechanical model. This representation in one dimension (i.e., if one of the random variables is binary) leads to a simple necessary and sufficient criterion for simulatability which is verifiable in deterministic time linear in $|\mathcal{X}| + |\mathcal{Y}| + |\mathcal{Z}|$. Moreover, the given algorithm yields the corresponding channel in case simulatability holds. In the general $n\,(\geq 2)$-dimensional case, an apparently close-to-tight (yet not sufficient in all cases) *necessary* criterion, which is checkable in time polynomial in $|\mathcal{X}| + |\mathcal{Y}| + |\mathcal{Z}|$ has been described. It is an open question however to find a simple necessary and sufficient criterion for the general case.

The introduced formalism can be helpful also with respect to other problems dealing with discrete distributions and noisy channels. An example is to determine the *intrinsic conditional information* $I(X;Y{\downarrow}Z)$ defined by

$$I(X;Y{\downarrow}Z) := \min_{XY \to Z \to \overline{Z}} I(X;Y|\overline{Z}) \ ,$$

a quantity that has been shown closely related to the possibility of secret-key agreement against *passive* adversaries [7], [10].

References

[1] R. Ahlswede and I. Csiszár, "Common randomness in information theory and cryptography - Part I: secret sharing," *IEEE Transactions on Information Theory*, vol. 39, no. 4, pp. 1121–1132, 1993.

[2] I. Csiszár and J. Körner, "Broadcast channels with confidential messages," *IEEE Transactions on Information Theory*, vol. IT-24, pp. 339–348, 1978.

[3] U. M. Maurer, "Information-theoretically secure secret-key agreement by NOT authenticated public discussion," *Advances in Cryptology - EU-*

ROCRYPT '97, Lecture Notes in Computer Science, vol. 1233, pp. 209–225, Springer-Verlag, 1997.

[4] U. M. Maurer, "The strong secret key rate of discrete random triples," *Communication and Cryptography – Two Sides of One Tapestry*, Kluwer Academic Publishers, pp. 271–285, 1994.

[5] U. M. Maurer, "Secret key agreement by public discussion from common information," *IEEE Transactions on Information Theory*, vol. 39, no. 3, pp. 733–742, 1993.

[6] U. M. Maurer and S. Wolf, "Secret-key agreement against active adversaries I: The robust secret-key rate," in preparation.

[7] U. M. Maurer and S. Wolf, "Unconditionally secure key agreement and the intrinsic conditional information," *IEEE Transactions on Information Theory*, vol. 45, no. 2, pp. 499–514, 1999.

[8] U. M. Maurer and S. Wolf, "Privacy amplification secure against active adversaries," *Advances in Cryptology - CRYPTO '97*, Lecture Notes in Computer Science, vol. 1294, pp. 307–321, Springer-Verlag, 1997.

[9] D. R. Stinson, "Universal hashing and authentication codes," *Advances in Cryptology - CRYPTO '91*, Lecture Notes in Computer Science, vol. 576, pp. 74–85, Springer-Verlag, 1992.

[10] S. Wolf, *Information-theoretically and computationally secure key agreement in cryptography*, ETH dissertation No. 13138, Swiss Federal Institute of Technology (ETH Zurich), May 1999.

[11] S. Wolf, "Strong security against active attacks in information-theoretic secret-key agreement," *Advances in Cryptology - ASIACRYPT '98*, Lecture Notes in Computer Science, vol. 1514, pp. 405–419, Springer-Verlag, 1998.

[12] A. D. Wyner, "The wire-tap channel," *Bell System Technical Journal*, vol. 54, no. 8, pp. 1355–1387, 1975.