Privacy for the paranoid ones

Artur Ekert & Alastair Kay

Among those who make their living out of science of secrecy, worrying and paranoia are just a sign of professionalism. Can we protect our secrets against those who wield superior technological powers? Here we study some bizarre correlations and offer some answers.

I. MISSION IMPOSSIBLE

Is there such a thing as a perfect cipher? If the history of secret communication is of any guidance here the answer is a resounding "no". But the situation is not as hopeless as it may sound. The requirements for perfect secrecy are well understood. All we need to construct a perfect cipher is shared randomness, more precisely, two identical, random and secret sequences of bits at two distant locations. Such sequences are called cryptographic keys. Any two parties who share a key, we call them Alice and Bob (not their real names, of course), can then use it to communicate secretly, employing a simple but nevertheless perfectly secure encryption method, known as the one-time pad. The message to be communicated is first written in binary and then encrypted by adding it, bit by bit, to the key. For example, given the key 1100101 Alice can encrypt her binary message, say 1011100, by combining each bit of the message with the respective bit of the key,

$$\begin{array}{r} 1\,0\,1\,1\,1\,0\,0 \\ +\,1\,1\,0\,0\,1\,0\,1 \\ \hline 0\,1\,1\,1\,0\,0\,1 \end{array}$$

The resulting cryptogram, 0111001, can be then publicly transmitted to Bob, who can recover the message by adding the cryptogram and the key. Both Alice and Bob must, of course, have exact copies of the key beforehand; Alice needs the key to encrypt the message, Bob needs the key to recover the message from the cryptogram. But for Eve, an eavesdropper who has intercepted the cryptogram and knows the general method of encryption but not the key, the message remains hidden. This is due to the randomness in the key. Since all possible keys are equally likely to be used the resulting cryptogram is completely independent from the message and no different from any randomly chosen binary sequence of the same length as the key and the message. The one-time-pad is therefore unbreakable, no matter what computational power the adversaries can muster.

There is a snag, however. All one-time pads suffer from a serious practical drawback, known as the key distribution problem. Each key can be used only once. In order to communicate secretly, Alice and Bob, who may be miles apart, must find a way to generate and distribute new keys. Moreover, they must be able to verify that the newly established keys are indeed secret, i.e. known only to them and nobody else. Solve this key distribution problem, combine it with the one-time pad, and you have constructed, behold, a truly unbreakable cipher!

Let us put all the practicalities aside, just for a moment, and dream about something that would solve the key distribution problem. For example, imagine that Alice and Bob found two magically linked coins, which always come out the same side up, either two heads or two tails, with equal probabilities. Alice and Bob can then toss such coins at their respective locations, writing 0 for heads and 1 for tails. The resulting binary strings will be random and identical, but will they be secret? Not necessarily. The same magic that correlates bits shared by Alice and Bob can correlate these bits to some other bits held by adversaries. This looks like a serious problem but, surprisingly, it can be fixed!

II. THE POWER OF FREE CHOICE

In order to achieve secrecy, we must let Alice and Bob do something that is beyond Eve's control. For example, Alice and Bob may be given a *choice* between two different coins; Alice can toss either coin A_1 or coin A_2 and Bob either B_1 or B_2. Please note, each of them can toss only one coin; tossing both A_1 and A_2 or B_1 and B_2 is not allowed. Suppose, again, the coins are magically linked; whenever A_1 and B_2 are tossed they always come out opposite but any other pair of coins always comes out the same. Other than that, heads and tails appear with equal probabilities. The magic can be succinctly summarised by the four conditions,

$$A_1 = B_1,\ B_1 = A_2,\ A_2 = B_2,\ \text{and}\ B_2 \neq A_1. \qquad (1)$$

The first, immediate, observation is that these conditions are contradictory. More precisely, it is impossible to assign values to A_1, A_2, B_1 and B_2, so that all the four conditions are satisfied. Let us stress, however, that this does not mean that the magic coins cannot exist. Remember, Alice and Bob can toss only one coin each thus they never test all the four conditions in one go, but only one of them at a time. This said, such correlations certainly preclude the existence of predetermined outcomes of coin tosses. Eve, who does not know beforehand which coins Alice and Bob are going to choose, cannot preprogram them without violating at least one of the four conditions. All pre-programmed coins can be detected by Alice and Bob for in any long sequence of tosses at least 25% of outcomes will deviate from the behaviour of the magic coins.

The second, more subtle, observation is that our magic correlations are "non-signalling". Alice, by making a

choice between A_1 and A_2, cannot signal one bit of information to Bob (and vice versa). The choice of the coin at one location does not affect the outcome of any coin toss at any other location. The table of probabilities helps to see that.

Alice		A_1		A_2	
Bob		0	1	0	1
B_1	0	$\frac{1}{2}$	0	$\frac{1}{2}$	0
	1	0	$\frac{1}{2}$	0	$\frac{1}{2}$
B_2	0	0	$\frac{1}{2}$	$\frac{1}{2}$	0
	1	$\frac{1}{2}$	0	0	$\frac{1}{2}$

The entries are the joint probabilities of outcomes given the choices of coins. For example, if Alice tosses coin A_1 and Bob B_2 then the probability that Alice gets 1 and Bob 0 is $\frac{1}{2}$. The table shows that the chances of getting 0 or 1 by one party (be it Alice or Bob) are not affected by the choices made by the other party. This property makes it impossible to extend the magic to any other coins. For suppose Bob, apart from B_1 and B_2, has another pair of coins, C_1 and C_2, which are also magically linked to Alice's coins, A_1 and A_2. Bob can then choose to toss both B_2 and C_1 and infer right away whether Alice tossed A_1 ($B_2 = C_1$) or A_2 ($B_2 \neq C_1$). This allows Alice to send instant messages to Bob. More than that, she is deprived of free will! Bob may toss B_2 and C_1 long before Alice makes her choice and if his outcomes tally Alice must choose A_1, otherwise she must choose A_2. All this means that the magic correlations must be *monogamous*, that is, they cannot be shared or extended beyond the two pairs of coins.

The two observations are crucial for what follows, so make sure your are comfortable with them before you read any further. The magic correlations and the sheer fact that Alice and Bob can now exercise their free will and choose which coin to toss turns the tables on Eve. She cannot pre-programme the coins and she cannot learn the outcomes by extending the magic correlations to another pair of coins for the magic correlations cannot be extended beyond the two legitimate parties. Whenever she tries the coins deviate from the magic correlations. Alice and Bob can detect this and the degree of deviation tells them to what extend someone tampered with the coins. All ingredients for a secure key distribution are now in place.

III. KEY DISTRIBUTION

Alice and Bob can establish a secret key by repeatedly tossing their magic coins and communicating in public. They discuss and compare the outcomes of selected coin tosses, and test them for evidence of eavesdropping.

For each toss Alice and Bob choose randomly, and independently from each other, which particular coin will be tossed: Alice is choosing between A_1 and A_2, Bob between B_1 and B_2. After the toss they announce publicly the type of the coin (but not the outcome of the toss!). The resulting sequence of data, at this stage, may look like this

Alice	A_1	A_2	A_1	A_1	A_2	A_1	A_1	A_1	A_2	A_1	A_1
	1	1	1	0	1	0	1	0	0	1	1
Bob	B_2	B_2	B_1	B_1	B_2	B_2	B_1	B_2	B_1	B_1	B_1
	0	1	1	0	1	1	1	1	0	1	1

where everything on the grey background is publicly known. For example, after the 42nd toss Alice may say, "I tossed A_1", and Bob, "And I tossed B_2" after which Alice records "Toss 42: A_1, outcome 0; B_2, outcome unknown" (for she does not know Bob's outcome) and Bob "Toss 42: A_1, outcome unknown; B_2, outcome 1" (for he does not know Alice's outcome). Here, we have tacitly assumed that Alice and Bob can communicate in public so that anybody, including Eve, can listen to, but nobody can alter these public messages. Think about a radio broadcasting or putting an advert in a newspaper.

Next Alice and Bob verify whether the coins are indeed magically linked. They perform a statistical check on some randomly selected records,

Alice	A_1	A_2	A_1	A_1	A_2	A_1	A_1	A_1	A_2	A_1	A_1
	1	1	1	0	1	0	1	0	0	1	1
Bob	B_2	B_2	B_1	B_1	B_2	B_2	B_1	B_2	B_1	B_1	B_1
	0	1	1	0	1	1	1	1	0	1	1

For example, Alice may announce, "In the 42nd toss I obtained 0", and Bob may answer, "And I got 1". The outcome is consistent with what they expect from the magic correlations so they tick off this record as "passed". If, after sufficiently many checks, Alice and Bob do not find any disagreement with the four conditions (**??**), they can be pretty confident that they are tossing untampered magic coins. Any deviation from the magic correlations indicates interference with the coins, possibly an attempt of eavesdropping, which forces Alice and Bob to abandon the protocol and start the key distribution from scratch.

However, once Alice and Bob are confident they are tossing untampered magic coins they know that the outcomes of the tosses that were not included in the statistical check and were not revealed in public are both secret and identical (unless Alice chose A_1 and Bob B_2, in which case the outcomes are always different and Bob can simply flip his bit value to make the outcomes identical). This way Alice and Bob end up with two identical sequences of bits each. In our simple example this will be 1100101. The key is established! Its secrecy based on the magic correlations and one innocuous but essential assumption — both Alice and Bob can *freely* choose which coins to toss. It seems that we have already achieved our goal. There is only one little problem with our, otherwise impeccable, solution of the key distribution problem, namely, the magic correlations do not exist! That is, we do not know of any physical process that can generate them. But stay calm, we can fix it.