Cryptography is the science that allows us to securely achieve various objectives while operating in an insecure environment. Let's try to picture some concrete scenarios where this may be desirable:

1. You are on the phone with your overseas friend and want to share a bit of gossip about your roommate, but you are worried that the roommate may be listening in on the conversation.
2. Your have a wireless router that you always use to access the internet, but your upstairs neighbor is also within the range of your router. You are worried about two things: (a) Your neighbor is snooping in on your internet traffic and (b) He is making illicit use of your paid internet connection without contributing.
3. You and your nine friends want to have a lucky draw for Chinese New Year – the winner gets a prize. However you are all in different places around the world, so you can only do it on the internet. Someone suggests to perform the draw at random, but you are all very suspicious that he will miraculously emerge as the "lucky" winner. Can you figure out a way to do the draw in a fair and judicious manner?
4. You are buying some candy from a guy on the street and he offers to let you pay with your Octopus card. But you don't know him and you are afraid that his card reader is a fake. Maybe it will transfer away all your credit to his account. Should you trust that this is not going to happen?
5. It is election day and you are heading to the polls to cast your vote. But you are planning to vote for an unpopular candidate, and if your boss (who is also running in the election) finds out you are afraid you'll get into trouble. Should you trust that your vote will not be revealed? How do you know it will be counted properly?

While these scenarios are very different, they all concern some objective to be achieved by legitimate parties via communication, but in a way that prevents malicious parties from gaining advantage by interfering with the communication in various ways. The objective is, by its nature, conflicting: It needs to be easy to achieve for legitimate users, but difficult to attain for the malicious ones. However, the participants are all people, or computers, so they all have access to essentially the same amount of resources. (If anything, it is the malicious users that are often willing to spend *more* resources – coming up with a successful scheme for stealing credit card numbers or uncovering the contents of diplomatic cables can make you instantly rich or famous.) How can you make the same task easy for one group of users, and hard for another one?

As it should already become apparent from these examples, there is wild variation in the tasks to be achieved, modes of communication, and adversarial roles. (For example, in scenario 1 it may be reasonable to assume that the adversary is a merely passive eavesdropper, while in scenario 4 a malicious user may be posing as a legitimate participant either from the buyer side or from the seller side). One highlight of cryptography is that is has been able to reduce these seemingly incomparable problems to a small number of *cryptographic primitives*, namely simple functionalities that, when used and combined properly, can solve a wide variety of problems in a secure manner.

What will I learn in this class (if I study)?

The simplest cryptographic tasks involves two participants, called Alice and Bob who are interacting over an insecure channel. An adversary, called Eve, is tampering with the channel in some way – she may be passively listening to the channel, or actively participating in the exchange. Alice wants to send some message to Bob in a way that will prevent Eve from "messing" with it. Here are two objectives that Alice and Bob may want to achieve:

- **Privacy:** Eve does not find out any information about the message.
- **Authenticity:** Bob is convinced that the message he received indeed originated from Alice and not from Eve.

To achieve either of these goals, Alice and Bob need to have advance knowledge of some secret called the *key* that Eve doesn't know. If, say, Bob doesn't have a key, then Alice and Bob cannot achieve privacy: Eve could impersonate Bob – namely, run the exact same algorithm that Bob does, and whatever Bob can learn by the end of the interaction, Eve can learn too.

In private-key cryptography, the key is a random string that Alice and Bob agree on in advance which is not known to Eve. Using this key, it is often possible to hide information from Eve which makes her task impossible, or at least very difficult. In all but the simplest scenarios, figuring out what to do with the key is far from obvious. One insight of cryptography is that unless we make some assumptions about the *computational power* of Alice, Bob, and Eve, any reasonable cryptographic functionality is impossible to achieve. However, if we are willing to make assumptions, then we can achieve lots of things.

In public-key cryptography, there is a pair of keys, a private key and a public key. The public key is known to everyone, including Eve, while the private key is known to only one of the participants. This allows for applications where Alice and Bob can communicate privately without having met in advance: Bob creates a pair of keys, keeps the private key to himself and publishes his public key. Alice can then (in certain settings and under certain assumptions) encode the message to Bob using his public key in a manner that is unintelligible to Eve, but perfectly understandable Bob.

Clearly public key cryptography is the more desirable alternative. So why do we bother to study the private key variant? One reason is pedagogical. The private key setting is simpler, so it will allow us to introduce our fist cryptographic concepts – definitions of security, the role of computational assumptions – without having to worry about other complications. A more important reason is that although public-key cryptography is more powerful, this power comes at a price. From a practical perspective, public-key schemes are more difficult to design, so they tend to be less efficient, and possibly easier to break. From a theoretical perspective, there is evidence that public-key cryptography requires us to make strictly stronger assumptions about the computational power of the adversary. We may want to avoid making such assumptions, unless it is absolutely necessary.

We will then see that these seemingly simple tasks can serve us as stepping stones for solving much more complicated problems. Preserving privacy and certifying authenticity of messages are very special cases of the problem of *secure mulitparty computation*: This is a general setting where a number of parties, some honest and some malicious, what to jointly perform some computational task in a way that keeps each participant's data private. We will see a general definition of secure multiparty computation and, time permitting, sketch how this task can be achieved. To enable this we will make use of a very powerful cryptographic device called a *zero-knowledge proof*: This

is a mathematical proof that convinces you a statement is true without revealing any information beyond the truth of the statement.

Cryptography is a vast subject with connections to other areas, and there is a number of exciting topics that we will (in all likelihood) not have enough time to address. You can consider studying (aspects of) these topics for your final project:

- **Composability** The models of cryptographic tasks we will consider in this class will be "closed" in the sense that they do not interact with the environment. However in the real world (say on the internet) it often happens that a single agent participates in several interactions at the same time. Can we guarantee that our protocols remain secure in such a setting?
- **The strength of assumptions** Virtually all cryptographic protocols require us to make some computational assumption about the strength of the adversary. A principal objective of cryptography is to *reduce* the security a of construction to a small number of well-studied assumptions. But some of these assumptions are stronger than others. In general we would like to know: What is the weakest possible assumption that enables a given cryptographic task?
- **Privacy** In cryptography, we want to preserve the secrecy of our data in a hostile environment. But we often want or need to engage in the conflicting task of *sharing* data, be it via email, facebook postings, or participation in the census. What kind of tradeoffs can we expect to achieve between the utility of sharing information and preserving its secrecy, and what is the best way to go about it?
- **Cryptography and game theory** In both cryptography and game theory, you have some agents participating in a protocol / game out of which they are trying to derive some benefit. Although the objectives are often different, tools from one area can sometimes help in the other.

The role of definitions

In cryptography, we always begin with a rigorous definition of the task we are trying to achieve. A cryptographic definition always consists of two parts: a *functionality requirement* and a *security requirement*. The functionality requirement describes the task that the legitimate parties want to achieve, while the security requirement specifies the kinds of things that we want to prevent the adversary from doing.

While definitions are arguably important in all fields of computer science, this is nowhere more so than in cryptography. For example, if we are looking at algorithms for maximum flow, the concept of "blocking flow" may be important because once we know it, the task of designing the algorithm becomes easier (maybe we try to greedily improve blocking flows) and the correctness of the algorithm is more transparent. However, even without knowing this concept, we would still be able to design and argue correctness of our algorithms, albeit in a more roundabout way.

In contrast, without definitions, we cannot even begin to think about cryptography. For instance consider scenario 1 above. We want to have some way of communicating over the phone securely, but what does "securely" even mean? It depends a lot of what the roommate is allowed to do.

Perhaps every time the roommate picks up the phone in the other room, you hear a little "click" and at that point you stop talking about her until she is out of the house. Perhaps there is no "click", but you can talk in French which your roommate doesn't speak. But maybe she will use a tape recorder and she will later ask someone to translate the conversation for her. Or – here is an even wilder scenario – maybe she made a recording of *you* when you last talked to your boyfriend and today she will replay that recording to impersonate you over the phone. Or maybe she'll point a gun to your head and ask you to follow her script in the conversation.

Then there is the question of what it is exactly that we are trying to prevent her from finding out. Certainly we do not want her to find out all the juicy details about the gossip, but maybe even merely saying her name over the phone will arouse her suspicion. Maybe even if she hears *nothing* about herself, she will find this unnatural and begin to suspect that she is being talked about behind her back at other times. How can we even begin to think about all these possibilities, and hundreds of other ones that we haven't even considered?

The point is that, unlike in most of computer science where we are only concerned with legitimate users of whatever we are designing, in cryptography we want our system to be secure against *all* "reasonable" behaviors of the adversary. But we cannot start designing such a system until we know what a reasonable behavior is.

A good definition can capture all these requirements in a succinct manner. This succinctness is important because it not only brings clarity of thought to the problem at hand, but it also allows us to reason rigorously about it. This reasoning often leads to surprising results. For example, an adversary that can affect the communication (let's call her adversary E) is clearly more powerful than one that can merely observe it (adversary F). So a definition that guarantees security against adversary A should intuitively be harder to achieve than one that only concerns adversary F. However, surprisingly often in cryptography it turns out that if you can achieve security against E, you can also achieve security against E! But unless we have a precise understanding of the power of E and F, there is no way to describe this phenomenon.

Another important aspect of cryptographic definitions is that they implicitly describe the limitations of the security model we have in mind, and of cryptographic solutions to security problems in general. This is especially important when we want to deploy cryptosystems in the real world. It could be that existing cryptographic solutions are too weak for the task at hand, yet stronger ones are impossible to achieve. In such cases system designers may need to resort to non-cryptoraphic solutions, such as economic incentives or legal measures. These kinds of issues are addressed by *system security*, which applies insights from cryptography but also other areas to address security issues in real-world systems.

1 Some attempts at encryption

Let us begin our study with the problem of *encryption*: Alice wants to send a message to Bob in a way that keeps the message secret from Eve. This is a task that people have been interested in way before the discovery of computers.

The Roman emperor Julius Caesar used the following cipher: Shift the position of every letter of the alphabet by three symbols forward. So `a` becomes `d`, `b` becomes `e`, and so on. So if he wanted to communicate this message to one of his generals:

```
meetmeattherubicon
```

(blank spaces omitted), instead he would write

```
phhwphdwwjhuxelfrq.
```

Needless to say Caesar's cipher can be easily broken.[1] In fact his scheme has a fundamental flaw which is unrelated to the specific choice of cipher: it is *deterministic*. A deterministic encryption scheme can never be useful: If Alice "encrypts" a message to Bob and Bob can decrypt it, then so can Eve because in the absence of "secret information" possessed by Bob, Eve can impersonate Bob in any interaction.

Here is a randomized variant of Caesar's scheme. If we identify the English alphabet symbols with the numbers 0 to 25, then Caesar's cipher is $x \to x+3 \mod 26$. To *decrypt*, one applies the inverse transformation $y \to y-3 \mod 26$. One thing Alice and Bob can do is to agree on a secret key $i \in \{0, \ldots, 25\}$ and apply the encoding $x \to x+i \mod 26$ instead.

This scheme is somewhat better than the previous one. If Eve now sees a *ciphertext* sent from Alice to Bob, she cannot decode it, because she does not know the key. However, she knows that the ciphertext could have come from at most 26 possible messages. She can now write out the list of these 26 messages. Chances are at most one of them will make any sense in English, in which case she has broken the scheme again.

However, notice that breaking the new cipher requires Eve to do 26 times the amount of work it took her to break the original one. This suggests that we could make the cipher harder to break by using a longer key. For example, instead of a one-letter key we could use a four-letter key like this:

```
    meetmeattherubicon
+   skltskltskltskltsk
----------------------
    eopmeolmlrpkmltvgx
```

Another scheme could work like this: To obtain the ciphertext, we still substitute symbols one by one, but instead of a cyclic shift, the public key is now an arbitrary permutation P of the English alphabet, like

```
x:    a b c d e f g h i j k l m n o p q r s t u v w x y z
P(x): l t r e w v s g m q u d f h i o b z x c p y k j a n
```

and the encoding is obtained by applying the transformation $x \to P(x)$ to each symbol separately.

In general, it seems reasonable that the longer you make the key, the harder it becomes to break the scheme; and in fact for short messages, these schemes may be difficult to break. However, when we try to use them to encrypt longer messages, the schemes become easier to break. In fact the message can be recovered without too much effort by analyzing the frequency of occurrences of the different symbols (or pairs or triplets of consecutive symbols) and matching them to the frequencies in standard English text.

[1] Perhaps it made *some* sense to use this cipher in Roman times: If most of the population was illiterate, yet they could recognize certain words by sight, shifting things a bit may have made it more difficult.

Even with short messages, the ciphertext could give us various clues about the message. For example, notice that the first sets of four characters in the ciphertext differ in only one place, which already tells us something about the message. Suppose the message is a list of directions to the hideaway of a terrorist, like

`leftrightleftleftrightleft`

Then it should not be hard to convince yourself that in any of these schemes, the message can be easily recovered. For more examples of historical encryption ciphers and ways to break them take a look at Chapter 1 of the textbook *Introduction to modern cryptography*.

Security by obscurity One implicit assumption that we made is that the encryption and decryption algorithms – the procedures used to encode and decode messages – are public and therefore known to the adversary. Shouldn't we be able to achieve better security by keeping the algorithms hidden from the adversary?

Most cryptographers believe that, in fact, the opposite is true: We can only gain more confidence in the security of the scheme by making its description public. In practice, this will invite more people to study the scheme and uncover its vulnerabilities. As a consequence, if after a long period of study no vulnerabilities are found, this should give us more confidence that the scheme is indeed secure. A historical example of a "secret" encryption scheme are the *enigma codes* used by the German military in the 1930s and during World War II. This scheme was implemented by an elaborate machine whose mode of operation was kept secret. It took years to break it, first by Polish codebreakers, and then (a later version) by a group in Britain including Alan Turing.

A more technical reason, at least in private-key cryptography, is that essentially nothing is gained by making the algorithms of the scheme secret, as the description of these algorithms can always be provided as part of the private key.

2 The one-time pad

As we observed, if we want to improve the security of encryption, it seems a good idea to use a longer key. Let's take this idea to an extreme. Suppose Alice wants to send Bob a message m consisting of k symbols. The *one-time pad* is the following encryption scheme: The private key is a random string r of k symbols. Alice sends the message

$$c = m + r = (m_1 + r_1, \ldots, m_k + r_k)$$

and to recover the message, given ciphertext c, Bob computes $c - r$.

This scheme should be very difficult to break: What Eve can observe is the string $c = m + r$, and because r is random, from Eve's point of view, the ciphertext is *statistically independent* of the message. This requirement on the distribution of the ciphertexts is known as *perfect security*; we will define it formally below.

The problem with this scheme is, of course, that it requires Alice and Bob to agree on a key that is as long as the message they are planning to exchange. This key is not reusable: If Alice wants to send a different message to Bob, she must use a different key. In effect, this requires Alice and

Bob to agree in advance on a key which is as long as the length of the messages they are ever going to exchange. This is impractical in many applications. Even if a method of distributing the keys can be arranged (say by physically shipping them on a memory stick), it may be hard to get hold of enough randomness to generate such long keys.

Can we get away with using shorter keys? If the objective is to have a ciphertext that is statistically independent of the message, then the answer is no. In short here is why. Let's assume (as we shall from now on) that the alphabet in which both messages and ciphertexts are written is $\{0,1\}$. Suppose Alice wants to send an n-bit message to Bob, but they have only exchanged a key of length k. So there are 2^n possible messages, and they can be encoded using 2^k different keys. So any ciphertext c could have arisen from at most 2^k different messages. Let m be one of these messages. Since $k < n$, there must be some message m' that does not encode to c under any key. So from the perspective of Eve, the ciphertexts of m and m' *are* distinguishable: While m can be encrypted into c with some probability (over the choice of key), m' can never yield ciphertext c.[2]

But perhaps we do not need to worry so much: Although *in principle* Eve can always distinguish between the encryptions of m and m', in practice it may be very difficult for her to do so. It may be that the only way to figure out if c can arise as the encryption of a message m is to go over all possible key and check if any of them works. This would take an amount of work proportional to 2^k, which even for reasonably small values of k, say $k = 128$, is forbidding.

In conclusion, even though we cannot achieve *perfect* security when the key is shorter than the message, there is still a glimmer of hope: Eve may not be able to break the encryption simply because it takes too much work for her to do so.

3 Definitions of encryption and security

Let us try to formalize the discussion on encryption based on our intuition about the one-time pad.

Recall that any cryptographic definition has two components: A *functionality requirement*, which essentially describes the problem we are trying to solve and the behavior of the honest parties – Alice and Bob, and a *security requirement*, which describes the security guarantees we want to obtain against the adversary.

In general, the functionality requirement is easier to state. Here we have Alice who wants to encrypt a message M using a key K. The functionality requirement says that when Alice and Bob share the same key, decryption does the opposite of encryption.

Definition 1. A *private-key encryption scheme* for message length m and key length k is a pair of algorithms (Enc, Dec) such that for every key K of length k and message M of length m, $Dec(K, Enc(K, M)) = M$.

One annoying thing about this definition is that it only works for a fixed message length and key length. What we really want is a single algorithm that we can apply to messages of arbitrary length, provided the key is long enough:

[2] This argument makes one unreasonably simplistic assumption: It applies to *deterministic* algorithms Enc and Dec, while in practice they can be randomized ones. It is not difficult to extend the argument so that it applies to randomized encryption and decryption algorithms as well.

Definition 2. A *private-key encryption scheme* is a pair of algorithms (Enc, Dec) and a function $m(k)$ such that for every integer k, every key K of length k and every message m of length $m(k)$, $Dec(K, Enc(K, M)) = M$.

Here k is called the *security parameter*: By making it larger, we hope to encrypt longer messages, and at the same time make the task of breaking the encryption harder for the adversary.

We now want to define security. Recall the kind of security we get from the one-time pad: For a random key, the distribution of the ciphertext is completely independent of the message that was encrypted. One way to say this is that for every pair of messages, the distributions of ciphertexts are identical.

Definition 3. A private-key encryption scheme (Enc, Dec) is *perfectly secure* if for every k and every pair of messages M, M' of length $m(k)$, the random variables $Enc(K, M)$ and $Enc(K, M')$ are identically distributed, where K is chosen uniformly at random from $\{0,1\}^k$.

As we discovered, if a private-key encryption scheme is perfectly secure, then $k \geq m(k)$. We will now consider ways to relax the definition of perfect security in a way that may allow us to bypass this restriction.

One thing we can do is relax the requirement that the two distributions are identical. Instead, we can merely require that they are close to one another. What does it mean for two distributions to be close? There are various measures of distance between distributions, but the one that is most relevant for us is *statistical distance*. Informally, two distributions are close in statistical distance if no observer can distinguish samples coming from one distribution and samples coming from the other one very often. More formally, we say the distributions of two random variables X and X' are ε-close in statistical distance if

$$|\Pr[A(X) = 1] - \Pr[A(X') = 1]| \leq \varepsilon$$

for every function A.

So perhaps instead of asking that $Enc(K, M)$ and $Enc(K, M')$ in the definition of perfect security are identically distributed, we can merely require that they are ε-close in statistical distance for some small ε. Unfortunately, this is not good enough: It turns out that as long as the key is shorter than the message, there will at least two messages M and M' so that the statistical distance between $Enc(K, M)$ and $Enc(K, M')$ is $1/2$ or more.