



CRYPTOGRAPHY
KNOWLEDGE AREA
(DRAFT FOR COMMENT)

AUTHOR: Nigel Smart – KU Leuven

EDITOR: George Danezis – University College London

REVIEWERS:

Dan Bogdanov - Cybernetica

Kenny Patterson – Royal Holloway, University of London

Liqun Chen – University of Surrey

Cryptography

Nigel P. Smart

April 2018

INTRODUCTION

The purpose of this chapter is to explain the various aspects of cryptography which we feel should be known to an expert in cyber-security. The presentation is at a level needed for an instructor in a module in cryptography; so they can select the depth needed in each topic. Whilst not all experts in cyber-security need be aware of all the technical aspects mentioned below, we feel they should be aware of all the overall topics and have an intuitive grasp as to what they mean, and what services they can provide. Our focus is mainly on primitives, schemes and protocols which are widely used, or which are suitably well studied that they could be used (or are currently being used) in specific application domains.

Cryptography by its very nature is one of the more mathematical aspects of cyber-security; thus this chapter contains a lot more mathematics than one has in some of the other chapters. The overall presentation assumes a basic knowledge of either first-year undergraduate mathematics, or that found in a discrete mathematics course of an undergraduate Computer Science degree.

The chapter is structured as follows: After a quick recap on some basic mathematical notation (Section 1), we then give an introduction to how security is defined in modern cryptography. This section (Section 2) forms the basis of our discussions in the other sections. Section 3 discusses information theoretic constructions, in particular the one-time pad, and secret sharing. Sections 4 and 5 then detail modern symmetric cryptography; by discussing primitives (such as block cipher constructions) and then specific schemes (such as modes-of-operation). Then in Sections 6 and 7 we discuss the standard methodologies for performing public key encryption and public key signatures, respectively. Then in Section 8 we discuss how these basic schemes are used in various standard protocols; such as for authentication and key agreement. All of the sections, up to and including Section 8, focus exclusively on constructions which have widespread deployment.

Section 9 begins our treatment of constructions and protocols which are less widely used; but which do have a number of niche applications. These sections are included to enable the instructor to prepare students for the wider applications of the cryptography that they may encounter as niche applications become more mainstream. In particular, Section 9 covers Oblivious Transfer, Zero-Knowledge, and Multi-Party Computation. Section 10 covers public key schemes with special properties, such as group signatures, identity-based encryption and homomorphic encryption.

The chapter assumes the reader wants to **use** cryptographic constructs in order to build secure systems, it is not meant to introduce the reader to attack techniques on cryptographic primitives. Indeed, all primitives here can be assumed to have been selected to avoid specific attack vectors, or key lengths chosen to avoid them. Thus, we do not discuss aspects in symmetric cryptography such as linear and differential cryptanalysis, or other forms of cryptanalysis, which would be needed by someone wanting to design their own primitives. On the public key side we therefore also do not discuss the large body of work on factoring, discrete logarithm algorithms, or on lattice reduction.

In a similar vein, we assume that all algorithms are implemented correctly, and with no side-channels. Again it should be assumed that any primitive given here which is implemented in a way which reveals

side-channel information to an attacker will necessarily be weaker. Much like cryptanalysis, how side-channel information is exploited depends on each specific primitive, and to some extent on the type of side-channel that is obtained.

CONTENT

1 Mathematics

[1, c8–c9, App B][2, c1–c5]

Cryptography is inherently mathematical in nature, and hence before proceeding we will set up some notation: The ring of integers is denoted by \mathbb{Z} , whilst the fields of rational, real and complex numbers are denoted by \mathbb{Q} , \mathbb{R} and \mathbb{C} . The ring of integers modulo N will be denoted by $\mathbb{Z}/N\mathbb{Z}$, when N is a prime p this is a finite field often denoted by \mathbb{F}_p . The set of invertible elements will be written $(\mathbb{Z}/N\mathbb{Z})^*$ or \mathbb{F}_p^* . An RSA modulus N will denote an integer N , which is the product of two (large) prime factors $N = p \cdot q$.

Finite abelian groups of prime order q are also a basic construct. These are either written multiplicatively, in which case an element is written as g^x for some $x \in \mathbb{Z}/q\mathbb{Z}$; when written additively an element can be written as $[x] \cdot P$. The element g (in the multiplicative case) and P (in the additive case) is called the generator.

The standard example of finite abelian groups of prime order used in cryptography are elliptic curves. An elliptic curve over a finite field \mathbb{F}_p is the set of solutions (X, Y) to an equation of the form

$$E : Y^2 = X^3 + A \cdot X + B$$

where A and B are fixed constants. Such a set of solutions, plus a special point at infinity denoted by \mathcal{O} , form a finite abelian group denoted by $E(\mathbb{F}_p)$. The group law is a classic law dating back to Newton and Fermat called the chord-tangent process. When A and B are selected carefully one can ensure that the size of $E(\mathbb{F}_p)$ is a prime q . This will be important later in Section 2.3 to ensure the discrete logarithm problem in the elliptic curve is hard.

Some cryptographic schemes make use of lattices which are discrete subgroups of the subgroups of \mathbb{R}^n . A lattice can be defined by a generating matrix $B \in \mathbb{R}^{n \times m}$, where each column of B forms a basis element. The lattice is then the set of elements of the form $\mathbf{y} = B \cdot \mathbf{x}$ where \mathbf{x} ranges over all elements in \mathbb{Z}^m . Since a lattice is discrete it has a well-defined length of the shortest non-zero vector. In Section 2.3 we note that finding this shortest non-zero vector is a hard computational problem.

Sampling a uniformly random element from a set A will be denoted by $x \leftarrow A$. If the set A consists of a single element a we will write this as the assignment $x \leftarrow a$; with the equality symbol $=$ being reserved for equalities as opposed to assignments. If A is a randomized algorithm, then we write $x \leftarrow A(y; r)$ for the assignment to x of the output of running A on input y with random coins r .

2 Cryptographic Security Models

[1, c1–c4][2, c11]

Modern cryptography has adopted a methodology of 'Provable Security' to define and understand the security of cryptographic constructions. The basic design procedure is to define the **syntax** for a cryptographic scheme. This gives the input and output behaviours of the algorithms making up the scheme and defines correctness. Then a **security model** is presented which defines what security goals are expected of the given scheme. Then, given a **specific instantiation** which meets the given syntax, a formal **security proof** for the instantiation is given relative to some known **hard problems**.

The security proof is not an absolute guarantee of security. It is a proof that the given instantiation, when implemented correctly, satisfies the given security model assuming some hard problems are

indeed hard. Thus, if an attacker can perform operations which are outside the model, or manages to break the underlying hard problem, then the proof is worthless. However, a security proof, with respect to well studied models and hard problems, can give strong guarantees that the given construction has no fundamental weaknesses.

In the next subsections we shall go into these ideas in more detail, and then give some examples of security statements.

2.1 Syntax of Basic Schemes

The syntax of a cryptographic scheme is defined by the algorithms which make up the scheme, as well as a correctness definition. The correctness definition gives what behaviour one can expect when there is no adversarial behaviour. For example, a symmetric encryption scheme is defined by three algorithms (KeyGen, Enc, Dec). The KeyGen algorithm is a probabilistic algorithm which outputs a symmetric key $k \leftarrow \text{KeyGen}()$; Enc is a probabilistic algorithm which takes a message $m \in \mathcal{M}$, some randomness $r \in \mathcal{R}$ and a key and returns a ciphertext $c \leftarrow \text{Enc}(m, k; r) \in \mathcal{C}$; whilst Dec is (usually) a deterministic algorithm which takes a ciphertext and a key and returns the underlying plaintext. The correctness definition is:

$$\forall k \leftarrow \text{KeyGen}(), r \leftarrow \mathcal{R}, m \leftarrow \mathcal{M}, \text{Dec}(\text{Enc}(m, k; r), k) = m.$$

For public key encryption schemes the definitions are similar, but now KeyGen() outputs key pairs and the correctness definition becomes:

$$\forall (pk, sk) \leftarrow \text{KeyGen}(), r \leftarrow \mathcal{R}, m \leftarrow \mathcal{M}, \text{Dec}(\text{Enc}(m, pk; r), sk) = m.$$

The equivalent constructions for authentication mechanisms are Message Authentication Codes (or MACs) in the symmetric key setting, and digital signatures schemes in the public key setting. A MAC scheme is given by a triple of algorithms (KeyGen, MAC, Verify), where the MAC function outputs a tag given a message and a key (and possibly some random coins), and the Verify function checks the message, tag and key are consistent. A signature scheme is given by a similar triple (KeyGen, Sign, Verify), where now the tag produced is called a 'signature'. Thus the correctness definitions for these constructions are as follows

$$k \leftarrow \text{KeyGen}(), r \leftarrow \mathcal{R}, m \leftarrow \mathcal{M}, \text{Verify}(m, \text{MAC}(m, k; r), k) = \text{true}.$$

and

$$(pk, sk) \leftarrow \text{KeyGen}(), r \leftarrow \mathcal{R}, m \leftarrow \mathcal{M}, \text{Verify}(m, \text{Sign}(m, sk; r), pk) = \text{true}.$$

Note, that for deterministic MACs the verification algorithm is usually just to recompute the MAC tag $\text{MAC}(m, k)$, and then check it was what was received.

2.2 Basic Security Definitions

A security definition is usually given in the context of an attacker's **security goal**, followed by their **capabilities**. So, for example, a naive security goal for encryption could be to recover the underlying plaintext, so-called One-Way (or OW) security. This process of an attacker trying to obtain a specific goal is called a *security game*, with the attacker winning the *game*, if they can break this security goal with greater probability than random guessing. This advantage in probability over random guessing is called the adversary's *advantage*. The capabilities are expressed in terms of what oracles, or functions, we give the adversary access to. So, for example, in a naive security game for encryption we may give the adversary no oracles at all, producing a so-called Passive Attack (or PASS) capability.

The attacker is modelled as an arbitrary algorithm, or Turing machine, A , and if we give the adversary access to oracles then we write these as subscripts $A^{\mathcal{O}}$. Thus in our naive security game, we define

the advantage of our adversary A , running in maximum time t , as:

$$\text{Adv}^{\text{OW-PASS}}(A, t) = \Pr \left[\begin{array}{l} \mathfrak{k} \leftarrow \text{KeyGen}(), \quad m^* \leftarrow \mathcal{M}, \quad r \leftarrow \mathcal{R}, \\ c^* \leftarrow \text{Enc}(m, \mathfrak{k}; r), \quad m \leftarrow A(c^*) : m = m^* \end{array} \right] - \frac{1}{|\mathcal{M}|}.$$

This assumes that the message space is not bigger than the space of all possible keys. We say that a given construction is secure in the given model (which our naive example would be named OW-PASS), if the above advantage is *negligible* for all probabilistic polynomial time adversaries A . Here, negligible and polynomial time are measured in terms of a *security parameter* (which one can think of as the key size). Note this is an asymptotic definition, which in the context of schemes with fixed key size, makes no sense. In such situations we require that (t/Adv) is greater than some given concrete bound such as 2^{128} , since it is believed that performing an algorithm requiring 2^{128} steps is infeasible even for a nation-state adversary.

In the context of encryption (both symmetric and public key) the above naive security goal is not seen as being suitable for real applications. Instead, the security goal of Indistinguishable encryptions (or IND) is usually used. This asks the adversary to first come up with two plaintexts, of equal length, and then the challenger (or environment) encrypts one of them and gives the resulting challenge ciphertext to the adversary. The adversary's goal is then to determine which plaintext was encrypted. In the context of a passive attack this gives an advantage statement of the form:

$$\text{Adv}^{\text{IND-PASS}}(A, t) = \Pr \left[\begin{array}{l} \mathfrak{k} \leftarrow \text{KeyGen}(), \quad b \leftarrow \{0, 1\}, \quad m_0, m_1, \text{state} \leftarrow A_1(), \\ r \leftarrow \mathcal{R}, \quad c^* \leftarrow \text{Enc}(m_b, \mathfrak{k}; r), \quad b' \leftarrow A_2(c^*, \text{state}) : b = b' \end{array} \right] - \frac{1}{2}.$$

where the two stages of the adversary are given by A_1 and A_2 .

In terms of encryption, the above passive attack is almost always not sufficient in terms of capturing real-world adversarial capabilities, since real systems almost always give the attacker additional attack vectors. Thus two other (increasingly strong) attack capabilities are usually considered. These are a Chosen Plaintext Attack (or CPA capability), in which the adversary is given access to an encryption oracle to encrypt arbitrary messages of his choice, and a Chosen Ciphertext Attack (or CCA capability), in which the adversary has both an encryption and decryption oracle. In the case of a public key scheme the adversary always has access to an encryption oracle because it can encrypt plaintexts for itself using the public key, so in this case PASS and CPA are equivalent. In the case of a CCA capability we restrict the decryption oracle so that the adversary may not ask of it the challenge ciphertext c^* ; otherwise it can trivially win the security game. Thus the advantage of an IND-CCA adversary against a public key encryption scheme would be defined as:

$$\text{Adv}^{\text{IND-CCA}}(A, t) = \Pr \left[\begin{array}{l} (\mathfrak{pk}, \mathfrak{sk}) \leftarrow \text{KeyGen}(), \quad b \leftarrow \{0, 1\}, \quad m_0, m_1, \text{state} \leftarrow A_1^{\text{Dec}(\cdot, \mathfrak{sk})}(\mathfrak{pk}), \\ r \leftarrow \mathcal{R}, \quad c^* \leftarrow \text{Enc}(m_b, \mathfrak{pk}; r), \quad b' \leftarrow A_2^{\text{Dec}(\cdot, \mathfrak{sk})}(c^*, \text{state}) : b = b' \end{array} \right] - \frac{1}{2}.$$

Other security definitions are possible for encryption (such as Real-or-Random) but the above are the main ones.

For MACs (resp. digital signature schemes) the standard security goal is to come up with a message/tag (resp. message/signature) pair which passes the verification algorithm, a so-called Universal Forgery (or UF) attack. We make no assumption about whether the message has any meaning, indeed, the attacker wins if he is able to create a signature on *any* bit-string. If the adversary is given no oracles then he is said to be mounting a passive attack, whilst if the adversary is given a tag generation (resp. signing oracle) he is said to be executing a Chosen Message Attack (CMA). In the latter case the final forgery must not be one of the outputs of the given oracle. In the case of MAC security,

one may also give the adversary access to a tag verification oracle. However, for deterministic MACs this is implied by the CMA capability and is hence usually dropped.

Again we define an advantage and require this to be negligible in the security parameter. For digital signatures the advantage for the UF-CMA game is given by:

$$\text{Adv}^{\text{UF-CMA}}(A, t) = \Pr \left[(\mathfrak{pk}, \mathfrak{sk}) \leftarrow \text{KeyGen}(), (m, \sigma) \leftarrow A^{\text{Sign}(\cdot, \mathfrak{sk})}(\mathfrak{pk}) : \text{Verify}(m, \sigma, \mathfrak{pk}) = \text{true} \right].$$

2.3 Hard Problems

As explained above, security proofs are always relative to some hard problems. These hard problems are often called *cryptographic primitives*, since they are the smallest atomic object from which cryptographic schemes and protocols can be built. Such cryptographic primitives come in two flavours: Either they are keyed complexity theoretic definitions of functions, or they are mathematical hard problems.

In the former case one could consider a function $F_k(\cdot) : D \rightarrow C$ selected from a function family $\{F_k\}$ and indexed by some index k (thought of as a key of varying length). One can then ask whether the function selected is indistinguishable (by a probabilistic polynomial time algorithm A which has oracle access to the function) from a uniform random function from D to C . If such an assumption holds, then we say the function family defines a (keyed) Pseudo-Random Function (PRF). In the case when the domain D is equal to the codomain C we can ask whether the function is indistinguishable from a randomly chosen permutation, in which case we say the family defines a (keyed) Pseudo-Random Permutation (PRP).

In the case of a block cipher, such as AES (see later), where one has $C = D = \{0, 1\}^{128}$, it is a basic assumption that the AES function family (indexed by the key k) is a Pseudo-Random Permutation.

In the case of mathematical hard problems we have a similar formulation, but the definitions are often more intuitive. For example, one can ask the question whether a given RSA modulus $N = p \cdot q$ can be factored into its prime components p and q , the so-called factoring problem. The RSA group $\mathbb{Z}/N\mathbb{Z}$ defines a finite abelian group of unknown order (the order is known to the person who created N), finding the order of this group is equivalent to factoring N . The RSA function $x \rightarrow x^e \pmod{N}$ is believed to be hard to invert, leading to the so-called RSA-inversion problem of, given $y \in (\mathbb{Z}/N\mathbb{Z})^*$, finding x such that $x^e = y \pmod{N}$. It is known that the function can easily be inverted if the modulus N can be factored, but it is unknown if inverting the function implies N can be factored. Thus we have a situation where one problem (factoring) seems to be harder to solve than another problem (the RSA problem). However, in practice, we assume that both problems are hard, given appropriately chosen parameters.

In finite abelian groups of known order (usually assumed to be prime), one can define other problems. The problem of inverting the function $x \rightarrow g^x$, is known as the Discrete Logarithm Problem (DLP). The problem of, given g^x and g^y , determining $g^{x \cdot y}$ is known as the Diffie–Hellman Problem (DHP). The problem of distinguishing between triples of the form (g^x, g^y, g^z) and $(g^x, g^y, g^{x \cdot y})$ for random x, y, z is known as the Decision Diffie–Hellman (DDH) problem. When written additively in an elliptic curve group, a DDH triple has the form $([x] \cdot P, [y] \cdot P, [z] \cdot P)$.

Generally speaking, the mathematical hard problems are used to establish the security of public key primitives. A major issue is that the above problems (Factoring, RSA-problem, DLP, DHP, DDH), on which we base all of our main existing public key algorithms, are easily solved by large-scale quantum computers. This has led designers to try to build cryptographic schemes on top of mathematical primitives which do not appear to be able to be broken by a quantum computer. Examples of such problems are the problem of determining the shortest vector in a high dimensional lattice, the so-called Shortest Vector Problem (SVP), and the problem of determining the closest lattice vector to a non-lattice vector, the so-called Closest Vector Problem (CVP). These two problems, and others, give rise to a whole new area called Post-Quantum Cryptography (PQC).

Example: Putting the above ideas together, one may encounter statements such as: *The public key encryption scheme XYZ is IND-CCA secure assuming the RSA-problem is hard and AES is a PRP.* This statement tells us that *any* attack against the XYZ scheme must either be against some weakness in the implementation, or must come from some attack not captured in the IND-CCA model, or must come from solving the RSA-problem, or must come from showing that AES is not a PRP.

2.4 Setup Assumptions

Some cryptographic protocols require some setup assumptions. These are assumptions about the environment, or some data, which need to be satisfied before the protocol can be considered secure. These assumptions come in a variety of flavours. For example, one common setup assumption is that there exists a so-called Public Key Infrastructure (or PKI), meaning that we have a trusted binding between entities' public keys and their identities.

Another setup assumption is the existence of a string (called the Common Reference String or *crs*) available to all parties, and which has been set up in a trusted manner, i.e. such that no party has control of this string.

Other setup assumptions could be physical, for example, that the algorithms have access to good sources of random numbers, or that their internal workings are not susceptible to an invasive attacker, i.e. they are immune to side-channel attacks.

2.5 Simulation and UC Security

The above definitions of security make extensive use of the notion of indistinguishability between two executions. Indeed, many of the proof techniques used in the security proofs construct *simulations* of cryptographic operations. A simulation is an execution which is indistinguishable from the real execution, but does not involve (typically) the use of any key material. Another method to produce security models is the so-called *simulation paradigm*, where we ask that an adversary cannot tell the simulation from a real execution (unless they can solve some hard problem). This paradigm is often used to establish security results for more complex cryptographic protocols.

A problem with both the game/advantage-based definitions defined earlier and the simulation definitions is that they only apply to stand-alone executions, i.e. executions of one instance of the protocol in one environment. To cope with arbitrarily complex executions and composition of cryptographic protocols an extension to the simulation paradigm exists called the Universal Composability (UC) framework.

3 Information-theoretically Secure Constructions

[1, c2][2, c19]

Whilst much of cryptography is focused against securing against adversaries which are modelled as probabilistic polynomial time Turing machines, some constructions are known to provide security against unbounded adversaries. These are called information-theoretically secure constructions.

3.1 One-Time Pad

The most famous primitive which provides information-theoretic security is the one-time pad. Here, a binary message $m \in \{0, 1\}^t$ is encrypted by taking a key $k \in \{0, 1\}^t$ uniformly at random, and then producing the ciphertext $c = m \oplus k$. In terms of our earlier security models, this is an IND-PASS scheme even in the presence of a computationally unbounded adversary. However, the fact that it does not provide IND-CPA security is obvious. The scheme is unsuitable in almost all modern environments as one requires a key as long as the message and the key may only be used once; hence the name one-time pad.

3.2 Secret Sharing

Secret sharing schemes allow a secret to be shared among a set of parties so that only a given subset can reconstruct the secret by bringing their shares together. The person who constructs the sharing of the secret is called the dealer. The set of parties who can reconstruct the secret are called qualified sets, with the set of all qualified sets being called an *access structure*.

Any set which is not qualified is said to be an unqualified set, and the set of all unqualified sets is called an *adversary structure*. The access structure is usually assumed to be monotone, in that if the parties in A can reconstruct the secret, then so can any super-set of A .

Many secret sharing schemes provided information-theoretic security, in that any set of parties which is unqualified can obtain no information about the shared secret even if they have unbounded computing power.

A special form of access structure is a so-called *threshold* structure. Here we allow any subset of $t + 1$ parties to reconstruct the secret, whereas any subset of t parties is unable to learn anything. The value t is being called the threshold. One example construction of a threshold secret sharing scheme for a secret s in a field \mathbb{F}_p , with $n > p$ is via Shamir's secret sharing scheme.

In Shamir secret sharing, one selects a polynomial $f(X) \in \mathbb{F}_p[X]$ of degree t with constant coefficients s , the value one wishes to share. The share values are then given by $s_i = f(i) \pmod{p}$, for $i = 1, \dots, n$, with party i being given s_i . Reconstruction of the value s from a subset of more than t values s_i can be done using Lagrange interpolation.

Due to an equivalence with Reed-Solomon error correcting codes, if $t < n/2$, then on receipt of n share values s_i , a reconstructing party can detect if any party has given it an invalid share. Additionally, if $t < n/3$ then the reconstructing party can correct for any invalid shares.

Replicated secret sharing is a second popular scheme which supports any monotone access structure. Given a boolean formula defining who should have access to the secret, one can define a secret sharing scheme from this formula by replacing all occurrences of AND with $+$ and all occurrences of OR with a new secret. For example, given the formulae

$$(P_1 \text{ AND } P_2) \text{ OR } (P_2 \text{ AND } P_3)$$

, one can share a secret s by writing it as $s = s_1 + s_2 = s'_2 + s_3$ and then, giving party P_1 the value s_1 , party P_2 the pair of values s_2 and s'_2 , and party P_3 the value s_3 . Replicated secret sharing is the scheme obtained in this way when putting the boolean formulae into Conjunctive Normal Form.

Of importance in applications of secret sharing, especially to Secure Multi-Party Computation (see Section 9.3) is whether the adversary structure is Q_2 or Q_3 . An adversary structure is said to be Q_i if no set of i unqualified sets have union the full set of players. Shamir's secret sharing scheme is Q_2 if $t < n/2$ and Q_3 when $t < n/3$. The error detection (resp. correction) properties of Shamir's secret sharing scheme mentioned above follow through to any Q_2 (resp. Q_3) adversary structure.

4 Symmetric Primitives

[1, c3–c6][2, c11–c14]

Symmetric primitives are a key component of many cryptographic constructions. There are three such basic primitives: block ciphers, stream ciphers, and hash functions. Theoretically, all are keyed functions, i.e. they take as input a secret key, whilst in practice one often considers hash functions which are unkeyed. At a basic level, all are functions $f : \mathcal{K} \times D \rightarrow C$ where \mathcal{K} is the key space, D is the domain (which is of a fixed finite size for block ciphers and stream ciphers).

4.1 Block Ciphers

A block cipher is a function $f : \mathcal{K} \times \{0, 1\}^b \rightarrow \{0, 1\}^b$, where b is the block size. Despite their names such functions should not be thought of as an encryption algorithm. It is, however, a building block in many encryption algorithms. The design of block ciphers is a deep area of subject in cryptography, analogous to the design of number theoretic one-way functions. Much like number-theoretic one-way functions, cryptographic constructions are proved secure relative to an associated hard problem which a given block cipher is assumed to satisfy.

For a fixed key, a block cipher is assumed to act as a permutation on the set $\{0, 1\}^b$, i.e. for a fixed key k , the map $f_k : \{0, 1\}^b \rightarrow \{0, 1\}^b$ is a bijection. It is also assumed that inverting this permutation is also easy (if the key is known). A block cipher is considered *secure* if no polynomial time adversary, given oracle access to a permutation on $\{0, 1\}^b$, can tell the difference between being given a uniformly random permutation or the function f_k for some fixed hidden key k , i.e. the block cipher is a PRP. In some applications, we only require that the block cipher is a function, i.e. not a bijection. In which case we require the block cipher is a PRF.

One can never prove that a block cipher is a PRP, so the design criteria is usually a task of building a mathematical construction which resists all known attacks. The main such attacks which one resists are so-called *linear cryptanalysis*, where one approximates non-linear components within the block cipher by linear functions, and *differential cryptanalysis*, where one looks at how two outputs vary on related input messages, e.g. one applies f_k to various inputs m_0 and m_1 where $m_0 \oplus m_1 = \Delta$ a fixed value.

The design of a block cipher is made up of a number of simpler components. There are usually layers of simple fixed permutations, and layers of table lookups. These table lookups are called S-boxes, where the S stands for substitutions. There are two main techniques to design block ciphers. Both repeat a simple operation (called a round) a number of times. Each round consists of a combination of permutations and substitutions, and a key addition. The main key is first expanded into round-keys, with each round having a different round-key.

In the first methodology, called a *Feistel Network*, the S-Boxes allowed in each round can be injective, i.e. non-invertible. Despite this, the Feistel constructions still maintain the overall invertibility of the block cipher construction. The second method is a *Substitution-Permutation Network* design in which each round consists of a round-key addition, followed by a fixed permutation, followed by the application of bijective S-boxes. In general, the Feistel construction requires more rounds than the Substitution-Permutation network construction.

The DES (Data Encryption Standard) block cipher (with an original key of 56-bits and block size of $b = 64$) is a Feistel construction. The DES algorithm dates from the 1970s, and the key size is now considered far too small for any application. However, one can extend DES into a 112- or 168-bit key block cipher to construct an algorithm called 2DES or 3DES. The use of 2DES or 3DES is still considered secure, although in some applications, the block size of 64-bits is considered insecure for real-world use.

The AES (Advanced Encryption Standard) block cipher is considered the modern replacement for DES, and it is a block cipher with a 128- or 256-bit key, and with a block size of $b = 128$ bits. The AES algorithm has hardware support on many microprocessors, making operations using AES much faster than using other cryptographic primitives.

4.2 Stream Ciphers

A stream cipher is one which produces an arbitrary length string of output bits, i.e. the codomain of the function is essentially unbounded. Stream ciphers can be constructed from block ciphers, by using a block cipher in Counter Mode (see Section 5.1). However, the stream cipher is usually

reserved for constructions which are special-purpose and for which the hardware complexity is much reduced.

Clearly, a stream cipher cannot be a permutation, but we require that no polynomial time adversary can distinguish oracle access to the stream cipher from oracle access to a uniformly random function with infinite codomain. The design of stream ciphers is more ad-hoc than that of the design of block ciphers. In addition, there is less widespread adoption outside specific application areas.

4.3 Hash Functions

Hash functions are much like block ciphers in that they should act as PRFs. However, the input domain can be unbounded. Since a PRF needs to be keyed to make any sense in theoretical tracts, a hash function is usually a keyed object. In practice, we often require an unkeyed object, in which case one considers the actual hash function used to have an implicit inbuilt fixed key, and have been chosen from a function family already.

When considering a fixed hash function, one is usually interested in the intractability of inverting the hash function (the one-way property), the intractability of finding two inputs with the same output (the collision resistance property), or the intractability of finding, given an input/output pair, a new input which gives the same output (the second-preimage resistance property).

Merkle-Damgård Construction: Early hash functions were based on the Merkle-Damgård construction. The family of such functions (MD4, MD5, SHA-1, SHA-256) have a number of issues, with only SHA-256 now being considered secure. The Merkle-Damgård construction takes a compression function f taking two input $f(x, y)$ of fixed length with the output length of x . This is used to derive a function which allows arbitrary length inputs by first dividing a message m into t blocks m_1, \dots, m_t each of length $|y|$, and then applying

$$h_i = f(h_{i-1}, m_i) \text{ for } i = 1, \dots, t,$$

where the output is h_t and h_0 is some initial value, which can be thought of as a fixed key for the hash function.

The above methodology requires a method to pad the initial input block to encode the length, and it suffers from a number of practical issues. For example, there are obvious length extension attacks which render the use of such hash functions problematic in some applications. For example, in HMAC (see Section 5.2), one requires two applications of the hash function to prevent such attacks.

Sponge Constructions: A more modern approach to hash function design is to create a so-called sponge construction. This is the design philosophy behind SHA-3 (a.k.a. Keccak). A sponge is a function which operates in two phases. In the first phase, one enters data into the sponge state and, in the second phase, one squeezes data out from the sponge.

The sponge state is a pair $(r, c) \in \{0, 1\}^{|r|+|c|}$, where the length of r denotes the input/output rate and c is a variable holding an internal state hidden from any attacker. The size of c is directly related to the security of the construction. At each round, a public permutation p is applied to the state (r, c) .

In the input phase of the sponge, the input data m , after suitable padding, is divided into t blocks m_1, \dots, m_t of size $|r|$. Then the state is updated via the mapping

$$(r_i, c_i) = p(r_{i-1} \oplus m_i, c_{i-1}) \text{ for } i = 1, \dots, t,$$

where r_0 and c_0 are initialized to be fixed all-zero bit strings. After data is entered into the sponge, one can obtain s blocks of $|r|$ -bit outputs o_1, \dots, o_s by computing

$$(o_i, c'_i) = p(o_{i-1}, c'_{i-1}) \text{ for } i = 1, \dots, s,$$

where $o_0 = r_t$ and $c'_0 = c_t$. Thus the whole function is given by

$$H(m_1, \dots, m_t) = o_1, \dots, o_s.$$

Random Oracle Model: Many cryptographic constructions are only secure if one assumes that the hash function used in the construction behaves 'like a random oracle'. Such constructions are believed to be secure in the real world, but theoretically, they are less pleasing. One can think of a proof of security in the random oracle model as a proof in which we allow the attacker to have their usual powers; however, when they (or any of the partners they are attacking) call the underlying hash function the call is made to an external party via an oracle call. This external party then simply plays back a random value, i.e. it does not use any algorithm to generate the random values. All that is required is that if the input is given to the oracle twice, then the same output is always returned.

This clearly does not capture attacks in which the adversary makes clever use of exactly how the hash function is defined etc, and how this definition interacts with other aspects of the scheme/protocol under analysis. However, this modelling methodology has proved remarkably good in enabling cryptographers to design schemes which are secure in the real world.

5 Symmetric Encryption and Authentication

[1, c3–c4][2, c13–c14]

A block cipher, such as AES or DES, does not provide an effective form of data encryption or data/entity authentication on its own. To provide such symmetric cryptographic constructions, one needs a scheme, which takes the primitive and then utilizes this in a more complex construction to provide the required cryptographic service. In the context of symmetric encryption, these are provided by modes of operation. In the case of authentication, it is provided by a MAC construction. Additionally, block ciphers are often used to take some entropy and then expand, or collapse, this into a pseudo-random stream or key; a so-called XOF (or Extendable Output Function) or KDF (or Key Derivation Function).

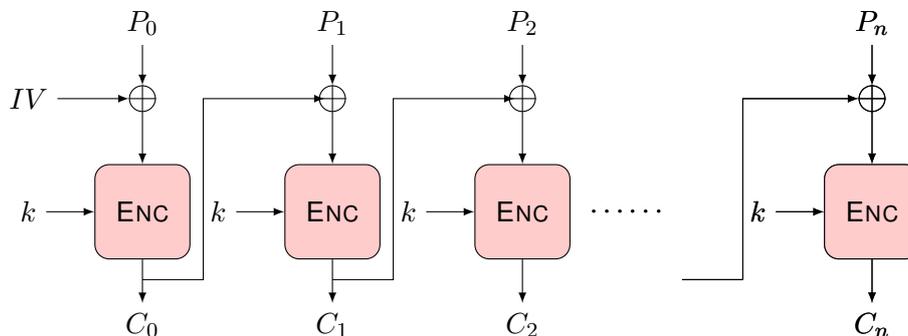


Figure 1: CBC Mode Encryption (All Figures are produced using *TikZ for Cryptographers* <https://www.iacr.org/authors/tikz/>).

5.1 Modes of Operation

Historically, there have been four traditional modes of operation to turn a block cipher into an encryption algorithm. These were ECB, CBC, OFB and CFB modes. In recent years, the CTR mode has also been added to this list. Among these, only CBC mode (given in Figure 1) and CTR mode (given in Figure 2) are used widely within current systems. In these Figures, the block cipher is represented by the function Enc

On their own, however, CBC and CTR modes only provide IND-CPA security. This is far weaker than the 'gold standard' of security, namely IND-CCA, thus, modern systems use modes which provide this level of security, also enabling additional data (such as session identifiers) to be tagged into the encryption algorithm. Such algorithms are called AEAD methods (or Authenticated Encryption with Authenticated Data). In such algorithms, the encryption primitive takes as input a message to be

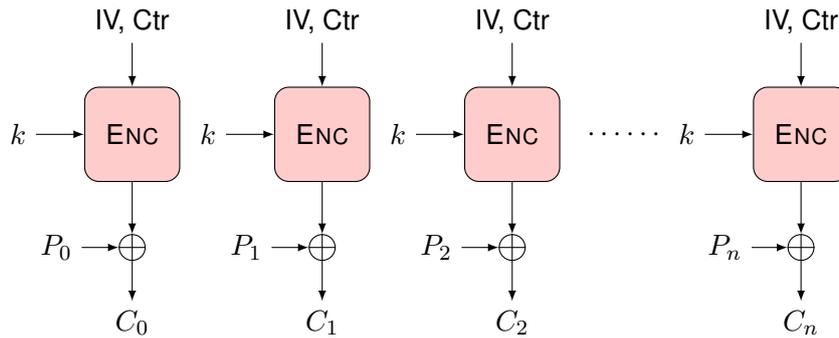


Figure 2: CTR Mode Encryption

encrypted, plus some associated data. To decrypt, the ciphertext is given, along with the associated data. Decryption will only work if both the key is correct and the associated data is what was input during the encryption process.

The simplest method to obtain an AEAD algorithm is to take an IND-CPA mode of operation such as CBC or CTR, and then to apply a MAC to the ciphertext and the data to be authenticated, giving us the so-called Encrypt-then-MAC paradigm. Thus, to encrypt m with authenticated data a , one applies the transform

$$c_1 \leftarrow \text{Enc}(m, \mathfrak{k}_1; r), \quad c_2 \leftarrow \text{MAC}(c_1 \| a, \mathfrak{k}_2; r),$$

with the ciphertext being (c_1, c_2) . In such a construction, it is important that the MAC is applied to the ciphertext as opposed to the message.

A major issue with the Encrypt-then-MAC construction is that one needs to pass the data to the underlying block cipher twice, with two different keys. Thus, new constructions of AEAD schemes have been given which are more efficient. The most widely deployed of these is GCM (or Galois Counter Mode), see Figure 3, which is widely deployed due to the support for this in modern processors.

One time AEAD constructions, otherwise known as DEMs, can be obtained by simply making the randomized AEAD deterministic by fixing the IV to zero.

5.2 Message Authentication Codes

Message authentication codes can be produced in roughly the same manner as modes of operation. In particular, the standard MAC function is to utilize CBC mode with a zero-IV, and then to output the final ciphertext block as the MAC tag, thus producing a deterministic MAC function. On its own, even with suitable padding of the message, this is only secure when used with fixed length messages. Thus, often a form of post-processing of the MAC output tag is performed. For example, the final CBC ciphertext block is then passed through another application of the underlying block cipher, but using a different key.

The GCM AEAD method of the previous section can be thought of as an Encrypt-then-MAC construction, with the IND-CPA encryption being CTR mode, and the MAC function being a function called GMAC. Although this is rarely used on its own as a MAC function.

Hash functions can also be used to construct MAC functions. The most famous of these is HMAC which is a construction designed for use with Merkle–Damgård-based hash functions. Since Merkle–Damgård-based hash functions suffer from length extension attacks, the HMAC function requires two applications of the underlying hash function. The construction produces a deterministic MAC function given by

$$\text{HMAC}(m, \mathfrak{k}) = H((\mathfrak{k} \oplus \text{opad}) \| H((\mathfrak{k} \oplus \text{ipad}) \| m)),$$

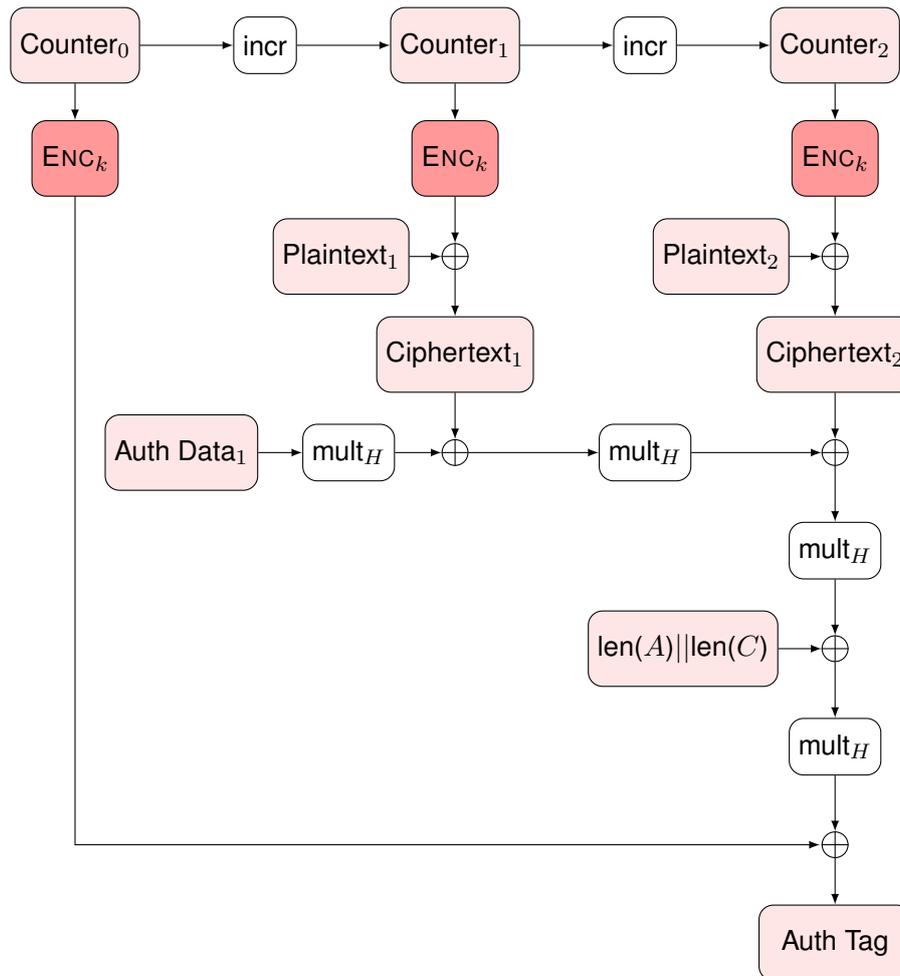


Figure 3: GCM Mode Encryption

where opad is the string 0x5c5c...5c5c and ipad is the string 0x3636...3636.

As HMAC is designed specifically for use with Merkle–Damgård-based hash functions, it makes no sense to use this construction when using a sponge based hash function such as SHA-3. The standardized MAC function derived from SHA-3 is called KMAC (or Keccak MAC). In this function, the sponge construction is used to input a suitably padded message, then the required MAC output is taken as the squeezed output of the sponge; whereas as many bits as squeezed are as needed for the MAC output.

5.3 Key Derivation and Extendable Output Functions

The security definition of a deterministic MAC is essentially equivalent to the definition that the output of the MAC function is indistinguishable from a random string, if one does not know the underlying secret key. As such, MAC functions can be used for other cryptographic operations. For example, in many situations, one must derive a long (or short) string of random bits, given some random input bits. Such functions are called KDFs or XOFs (for Key Derivation Function and Extendable Output Function). Usually, one uses the term KDF when the output is of a fixed length, and XOF when the output could be of an arbitrary length. But the constructions are, usually, essentially the same in both cases.

Such functions can take an arbitrary length input string, and produce another arbitrary length output string which is pseudo-random. There are three main constructions for such functions; one based on

block ciphers, one on the Merkle–Damgård hash functions, and one based on sponge-based hash functions.

The constructions based on a block cipher are, at their heart, using CBC-MAC, with a zero key to compress the input string into a cryptographic key and then use the CTR mode of operation under this key to produce the output string. Hence, the construction is essentially given by

$$\mathfrak{k} \leftarrow \text{CBC-MAC}(m, \mathbf{0}), \quad o_1 \leftarrow \text{Enc}(\mathbf{1}, \mathfrak{k}), \quad o_2 \leftarrow \text{Enc}(\mathbf{2}, \mathfrak{k}), \quad \dots$$

where Enc is the underlying block cipher.

The constructions based on the Merkle–Damgård hash function use a similar structure, but using one hash function application per output block, in a method similar to the following

$$o_1 \leftarrow H(m\|1), \quad o_2 \leftarrow H(m\|2), \quad \dots$$

Due to the way Merkle–Damgård hash functions are constructed, the above construction (for large enough m) can be done more efficiently than simply applying H as many times as the number of output blocks will dictate.

As one can imagine, the functions based on Keccak are simpler—one simply inputs the suitably padded message into the sponge and then squeezes as many output bits out as required.

6 Public Key Encryption

[1, c11][2, c15–c17]

As explained above, public key encryption involves two keys, a public one \mathfrak{pk} and a private one \mathfrak{sk} . The encryption algorithm uses the public key, whilst the decryption algorithm uses the secret key. The standard security requirement for public key encryption is that the scheme should be IND-CCA. Note that since the encryption key is public we have that IND-PASS is the same as IND-CPA for a public key encryption scheme.

6.1 KEM-DEM Philosophy

In general, public key encryption schemes are orders of magnitude less efficient than symmetric key encryption schemes. Thus, the usual method in utilizing a public key scheme, when large messages need to be encrypted, is via a hybrid method. This hybrid methodology is called the KEM-DEM philosophy. A KEM, which stands for Key Encapsulation Mechanism, is a public key method to transmit a short key, selected at random from a set \mathcal{K} , to a designated recipient. Whereas, a DEM, or Data Encryption Mechanism, is essentially the same as an IND-CCA symmetric encryption scheme, which has key space \mathcal{K} . Since a DEM is only ever used once with the same key, we can actually use a weaker notion of IND-CCA encryption for the DEM, in which the adversary is not given access to an encryption oracle; which means the DEM can be *deterministic*.

For a KEM, we call the encryption and decryption mechanisms encapsulation and decapsulation, respectively. It is usual for the syntax of the encapsulation algorithm to not take any input, bar the randomness, and then to return both the ciphertext and the key which it encapsulates. Thus, the syntax, and correctness, of a KEM becomes

$$(\mathfrak{pk}, \mathfrak{sk}) \leftarrow \text{KEMKeyGen}(), \quad r \leftarrow \mathcal{R}, \quad (\mathfrak{k}, c) \leftarrow \text{KEMEnc}(\mathfrak{pk}; r), \quad \text{KEMDec}(c, \mathfrak{sk}) = \mathfrak{k}.$$

The security definition for a KEM becomes

$$\text{Adv}^{\text{IND-CCA}}(A, t) = \Pr \left[(\mathfrak{pk}, \mathfrak{sk}) \leftarrow \text{KEMKeyGen}(), \quad b \leftarrow \{0, 1\}, \quad \mathfrak{k}_0 \leftarrow \mathcal{K}, \quad r \leftarrow \mathcal{R}, \right. \\ \left. \mathfrak{k}_1, c^* \leftarrow \text{KEMEnc}(\mathfrak{pk}; r), \quad b' \leftarrow A_2^{\text{KEMDec}(\cdot, \mathfrak{sk})}(c^*, \mathfrak{k}_b) : b = b' \right] - \frac{1}{2}.$$

To construct the hybrid public key encryption scheme we define $\text{KeyGen}()$ to be equal to KEMKeyGen , then $\text{Enc}(m, \text{pk}; r)$ outputs (c_0, c_1) where

$$\text{k}, c_0 \leftarrow \text{KEMEnc}(\text{pk}; r), \quad c_1 \leftarrow \text{DEM}(m, \text{k}),$$

with $\text{Dec}((c_0, c_1), \text{sk})$ being given by

$$\text{k} \leftarrow \text{KEMDec}(c_0, \text{sk}), \quad m \leftarrow \text{DEM}^{-1}(c_1, \text{k}).$$

6.2 Constructions based on RSA

The simplest public key encryption scheme is the RSA scheme, which is based on the difficulty of factoring integers. In the key generation algorithm, two large primes p and q are selected and multiplied together to form $N = p \cdot q$. Then a (usually small) integer e is selected which is co-prime to $\phi(N) = (p - 1) \cdot (q - 1)$. Finally, the integer d is found, using the extended Euclidean algorithm, such that $d = 1/e \pmod{\phi(N)}$. The public key is set to be $\text{pk} = (N, e)$, whereas the secret key is set to be $\text{sk} = (N, d)$. Note that given the pk only, finding the secret key sk is provably equivalent to factoring N .

The public/private keys are used via the RSA function $x \rightarrow x^e \pmod{N}$, which has the inverse map $x \rightarrow x^d \pmod{N}$. Thus, the RSA function is a trapdoor permutation on the group $(\mathbb{Z}/N\mathbb{Z})^*$. It is not believed that inverting the RSA map is equivalent to factoring, so inversion of this map is identified as a separate problem (the RSA problem). At the time of writing, one should select p and q to be primes of at least 2048 bits in length to obtain suitable security.

There are many historic ways of using the RSA function as a public key encryption scheme, many of which are now considered obsolete and/or insecure. The two recommended methodologies in using RSA for encryption are RSA-OAEP and RSA-KEM; which are both IND-CCA secure in the random oracle model.

OAEP, or Optimized Asymmetric Encryption Padding, is a method to use the RSA function directly as an IND-CCA public key encryption algorithm. OAEP is parametrized by two integers $k_0, k_1 \geq 128$ such that $n = \log_2 N - k_0 - k_1 \geq 0$. We then encrypt messages of at most n bits in length as follows, using hash-functions (which are assumed to be random oracles for the security proof)

$$\begin{aligned} G : \{0, 1\}^{k_0} &\rightarrow \{0, 1\}^{n+k_1} \\ H : \{0, 1\}^{n+k_1} &\rightarrow \{0, 1\}^{k_0}. \end{aligned}$$

We then encrypt using the function

$$c \leftarrow \left(\left\{ \left(m \parallel 0^{k_1} \right) \oplus G(R) \right\} \parallel \left\{ R \oplus H \left(\left(m \parallel 0^{k_1} \right) \oplus G(R) \right) \right\} \right)^e \pmod{N}$$

where

- $m \parallel 0^{k_1}$ means m followed by k_1 zero bits,
- R is a random bit string of length k_0 ,
- \parallel denotes concatenation.

RSA-KEM, on the other hand, is a KEM which is much simpler to execute. To produce the encapsulated key and the ciphertext, one takes the random input r (which one thinks of as a uniformly random element in $(\mathbb{Z}/N\mathbb{Z})^*$). Then the KEM is defined by

$$c \leftarrow r^e \pmod{N}, \quad \text{k} \leftarrow H(r),$$

where $H : (\mathbb{Z}/N\mathbb{Z}) \rightarrow \mathcal{K}$ is a hash function, which we model as a random oracle.

6.3 Constructions based on Elliptic Curves

Elliptic Curve Cryptography, or ECC, uses the fact that elliptic curves form a finite abelian group. In terms of encryption schemes, the standard method is to use ECIES (Elliptic Curve Integrated Encryption Scheme) to define a public key, KEM which is IND-CCA in the random oracle model, assuming the DDH problem in the subgroup of the elliptic curve being used. In practice, this means that one selects a curve $E(\mathbb{F}_p)$ for which there is a point $P \in E(\mathbb{F}_p)$ whose order is a prime $q > 2^{256}$.

For ECIES, the KeyGen algorithm is defined as follows. A secret key $\mathfrak{sk} \leftarrow \mathbb{F}_q^*$ is selected uniformly at random, and then the public key is set to be $Q \leftarrow [\mathfrak{sk}]P$. Key encapsulation is very similar to RSA-KEM in that it is defined by

$$r \leftarrow \mathbb{F}_q^*, C \leftarrow [r] \cdot P, \mathfrak{k} \leftarrow H([r] \cdot Q),$$

where $H : E(\mathbb{F}_p) \rightarrow \mathcal{K}$ is a hash function (modelled as a random oracle). To decapsulate the key is recovered via

$$\mathfrak{k} \leftarrow H([\mathfrak{sk}]C).$$

Compared to RSA-based primitives, ECC-based primitives are relatively fast and use less bandwidth. This is because, at the time of writing, one can select elliptic curve parameters with $p \approx q \approx 2^{256}$ to obtain security equivalent to a work-factor of 2^{128} operations. Hence, in current systems elliptic curve-based systems are preferred over RSA-based ones.

6.4 Lattice-based Constructions

A major problem with both RSA and ECC primitives is that they are not secure against quantum computers; namely, Shor's algorithm will break both the RSA and ECC hard problems in polynomial time. Hence, the search is on for public key schemes which would resist the advent of a quantum computer. The National Institute of Standards and Technology (NIST) is currently engaged in a process to determine potential schemes which are *post-quantum* secure.

The most prominent of these so-called post-quantum schemes are those based on hard problems on lattices. In particular, the NTRU schemes and a variety of schemes based on the Learning With Errors (LWE) problem, and its generalisation to polynomial rings, known as the Ring-LWE problem.

7 Public Key Signatures

[1, c12][2, c16]

Just as with public key encryption algorithms, modern digital signature algorithms are based either on the RSA problem or a variant of the discrete logarithm problem. For post-quantum security, there are a number of proposals based on lattice constructions; but none have yet been widely accepted or deployed at the time of writing this document.

The prototypical digital signature scheme given in test-books is loosely called RSA-FDH, where FDH stands for Full Domain Hash. The algorithm takes a message m and signs it by outputting

$$s = H(m)^d \pmod{N}.$$

Verification is then performed by testing whether the following equation holds

$$s^e = H(m) \pmod{N}.$$

Here, the hash function is assumed to have domain $\{0, 1\}^*$ and co-domain $(\mathbb{Z}/N/\mathbb{Z})^*$. This scheme comes with a security proof in the random oracle model, but is almost impossible to implement as no standardized hash function has co-domain the whole of $(\mathbb{Z}/N/\mathbb{Z})^*$, since N is much bigger than the output of hash functions such as SHA-256.

All standardized hash functions output a value in $\{0, 1\}^t$ for some t , thus what is usually done is to take a hash value and then prepend it with some known pre-determined values, and then 'sign' the result. This forms the basic idea behind the Public Key Cryptography Standards (PKCS) v1.5 signature standard. This signs a message by computing

$$s = (0x01\|0xFF \dots 0xFF\|0x00\|H(m))^d \pmod{N},$$

where enough padding of $0xFF$ bytes is done to ensure the whole padded string is just less than N in size. Despite the close relationship to RSA-FDH, the above signature scheme has no proof of security, and hence a more modern scheme is usually to be preferred.

7.1 RSA-PSS

The modern way to use the RSA primitive in a digital signature scheme is via the padding method called PSS (Probabilistic Signature Scheme). This is defined much like RSA-OAEP via the use of two hash functions, one which expands data and one which compresses data:

$$\begin{aligned} G : \{0, 1\}^{k_1} &\longrightarrow \{0, 1\}^{k-k_1-1}, \\ H : \{0, 1\}^* &\longrightarrow \{0, 1\}^{k_1}, \end{aligned}$$

where $k = \log_2 N$. From G we define two auxiliary functions

$$G_1 : \{0, 1\}^{k_1} \longrightarrow \{0, 1\}^{k_0}$$

which returns the first k_0 bits of $G(w)$ for $w \in \{0, 1\}^{k_1}$,

$$G_2 : \{0, 1\}^{k_1} \longrightarrow \{0, 1\}^{k-k_0-k_1-1}$$

which returns the last $k - k_0 - k_1 - 1$ bits of $G(w)$ for $w \in \{0, 1\}^{k_1}$, i.e. $G(w) = G_1(w)\|G_2(w)$.

To sign a message m the private key holder performs the following steps:

- $r \leftarrow \{0, 1\}^{k_0}$.
- $w \leftarrow H(m\|r)$.
- $y \leftarrow 0\|w\|(G_1(w) \oplus r)\|G_2(w)$.
- $s \leftarrow y^d \pmod{N}$.

To verify a signature (s, m) the public key holder performs the following

- $y \leftarrow s^e \pmod{N}$.
- Split y into the components $b\|w\|\alpha\|\gamma$ where b is one bit long, w is k_1 bits long, α is k_0 bits long and γ is $k - k_0 - k_1 - 1$ bits long.
- $r \leftarrow \alpha \oplus G_1(w)$.
- The signature is verified as correct if and only if b is the zero bit, $G_2(w) = \gamma$ and $H(m\|r) = w$.

Despite being more complicated than PKCS-1.5, the RSA-PSS scheme has a number of advantages. It is a randomized signature scheme, i.e. each application of the signing algorithm on the same message will produce a distinct signature, and it has a proof of security in the random oracle model relative to the difficulty of solving the RSA problem.

7.2 DSA, EC-DSA and Schnorr Signatures

The standard methodology for performing signatures in the discrete logarithm setting is to adapt interactive verification protocols, using proofs of knowledge of a discrete logarithm (see Sections 8.1 and 9.2), and then to convert them into a non-interactive signature scheme using a hash function.

The two most well known of these are the DSA (Digital Signature Algorithm) method and a method due to Schnorr. The former has widespread deployment but establishing security via security proofs uses less well-accepted assumptions, whereas the latter is less deployed but has well-established security proofs. The former also has, as we shall see, a more complex signing process. Both cases use a hash function H with codomain $(\mathbb{Z}/q\mathbb{Z})^*$, unlike RSA-FDH this is easy to construct as q is relatively small.

We will describe both algorithms in the context of elliptic curves. Both make use of a public key of the form $Q = [x] \cdot P$, where x is the secret key. To sign a message m , in both algorithms, one first selects a random value $k \in (\mathbb{Z}/q\mathbb{Z})^*$, and computes $r \leftarrow x - \text{coord}([k] \cdot P)$. One then computes a hash of the message. In the DSA algorithm, this is done with $e \leftarrow H(m)$, whereas for the Schnorr algorithm, one computes it via $e \leftarrow H(m||r)$. Then the signature equation is applied which, in the case of EC-DSA, is

$$s \leftarrow (e + x \cdot r)/k \pmod{q}$$

and, in the case of Schnorr, is

$$s \leftarrow (k + e \cdot x) \pmod{q}.$$

Finally, the output signature is given by (r, s) for EC-DSA and (e, s) for Schnorr.

Verification is done by checking the equation

$$r = x - \text{coord}([e/s] \cdot P + [r/s] \cdot Q)$$

in the case of EC-DSA, and by checking

$$e = H(m||x - \text{coord}([s] \cdot P - e \cdot Q))$$

in the case of Schnorr. The key difference in the two algorithms is not the signing and verification equations (although these do affect performance), but the fact that, with the Schnorr scheme, the r value is also entered into the hash function to produce e . This small distinction results in the different provable security properties of the two algorithms.

A key aspect of both EC-DSA and Schnorr signatures is that they are very brittle to exposure of the per-message random nonce k . If only a small number of bits of k leak to the attacker with every signing operation, then the attacker can easily recover the secret key.

8 Standard Protocols

[2, c18]

Cryptographic protocols are interactive operations conducted between two or more parties in order to realize some cryptographic goal. Almost all cryptographic protocols make use of the primitives we have already discussed (encryption, message authentication, secret sharing). In this section, we discuss the two most basic forms of protocol, namely authentication and key agreement.

8.1 Authentication Protocols

In an authentication protocol, one entity (the Prover) convinces the other entity (the Verifier) that they are who they claim to be, and that they are 'online'; where 'online' means that the verifying party is assured that the proving party is actually responding and it is not a replay. There are three basic types of protocol: Encryption based, Message Authentication based and Zero-Knowledge based.

Encryption-Based Protocols: These can operate in the symmetric or public key setting. In the symmetric key setting, both the prover and the verifier hold the same secret key, whilst in the public key setting, the prover holds the private key and the verifier holds the public key. In both settings, the verifier first encrypts a random nonce to the prover, the prover then decrypts this and returns it to the verifier, the verifier checks that the random nonce and the returned value are equivalent.

$$\begin{array}{ccc}
 \text{Verifier} & & \text{Prover} \\
 N \leftarrow \mathcal{M} & & \\
 c \leftarrow \text{Enc}(N, \text{pk}; r) & \xrightarrow{c} & \\
 & \xleftarrow{m} & m \leftarrow \text{Dec}(c, \text{sk}) \\
 N \stackrel{?}{=} m & &
 \end{array}$$

The encryption scheme needs to be IND-CCA secure for the above protocol to be secure against active attacks. The nonce N is used to prevent replay attacks.

Message Authentication-Based Protocols: These also operate in the public key or the symmetric setting. In these protocols, the verifier sends a nonce in the clear to the prover, the prover then produces a digital signature (or a MAC in the symmetric key setting) on this nonce and passes it back to the verifier. The verifier then verifies the digital signature (or verifies the MAC).

$$\begin{array}{ccc}
 \text{Verifier} & & \text{Prover} \\
 N \leftarrow \mathcal{M} & \xrightarrow{N} & \\
 & \xrightarrow{\sigma} & \sigma \leftarrow \text{Sign}(N, \text{sk}) \\
 \text{Verify}(N, \sigma, \text{pk}) \stackrel{?}{=} \text{true} & &
 \end{array}$$

Zero-Knowledge-Based Zero-knowledge-based authentication protocols are the simplest examples of zero-knowledge protocols (see Section 9.2) available. The basic protocol is a so-called Σ - (or Sigma-) protocol consisting of three message flows; a commitment, a challenge and a response. The simplest example is the Schnorr identification protocol, based on the hardness of computing discrete logarithms. In this protocol, the Prover is assumed to have a long-term secret x and an associated public key $Q = [x] \cdot P$. One should note the similarity of this protocol to the Schnorr signature scheme above.

$$\begin{array}{ccc}
 \text{Verifier} & & \text{Prover} \\
 & \xleftarrow{r} & k \leftarrow \mathbb{Z}/q\mathbb{Z} \\
 & & R \leftarrow [k] \cdot P \\
 e \leftarrow \mathbb{Z}/q\mathbb{Z} & \xrightarrow{e} & \\
 & \xrightarrow{s} & s \leftarrow (k + e \cdot x) \pmod{q} \\
 R \stackrel{?}{=} [s] \cdot P - e \cdot Q & &
 \end{array}$$

Indeed, the conversion of the Schnorr authentication protocol into the Schnorr signature scheme is an example of the Fiat–Shamir transform, which transforms any Σ -protocol into a signature scheme. If the underlying Σ -protocol is secure, in the sense of a zero-knowledge proofs of knowledge (see Section 9.2), then the resulting signature scheme is UF-CMA.

8.2 Key Agreement Protocols

A key agreement protocol allows two parties to agree on a secret key for use in subsequent protocols. The security requirements of key agreement protocols are very subtle, leading to various subtle security properties that many deployed protocols may or may not have. The basic security requirements are

- The underlying key should be indistinguishable from random to the adversary, or that at least it should be able to be used in the subsequent protocol without the adversary breaking the subsequent protocol.

- Each party is assured that only the other party has access to the secret key. This is so-called mutual authentication. In many application scenarios (e.g. in the Transport Layer Security (TLS) protocol), one only requires this property of one-party, in which case we are said to only have one-way authentication.

Kerberos is an example of a (usually) symmetric key-based key agreement system. This is a protocol that requires trusted parties to relay and generate secret keys from one party to another. It is most suited to closed corporate networks. On the public internet, protocols like Kerberos are less useful. Thus, here one uses public key-based protocols such as TLS and IPSec. More advanced properties required of modern public key-based protocols are as follows.

- **Key Confirmation:** The parties know that the other party has received the same secret key. Sometimes this can be eliminated as the correct execution of the subsequent protocol using the secret key provides this confirmation. This later process is called *implicit key confirmation*.
- **Forward Secrecy:** The compromise of a participants long-term secret in the future does not compromise the security of the secret key derived now, i.e. current conversations are still secure in the future.
- **Unknown Key Share Security:** This prevents one party (Alice) sharing a key with Bob, whereas Bob thinks he shares a key with Charlie, despite sharing it with Alice.

Key Transport: The most basic form of key agreement protocol is a form of key transport in which the parties use public key encryption to exchange a random key. In the case of a one-way authenticated protocol, this was the traditional method of TLS operation (up until TLS version 1.2) between a server and a client

$$\begin{array}{ccc}
 \text{Client} & & \text{Server} \\
 & \xleftarrow{pk} & \\
 k \leftarrow \mathcal{K} & & \\
 c \leftarrow \text{Enc}(k, pk; r) & \xrightarrow{c} & k \leftarrow \text{Enc}(c, sk)
 \end{array}$$

This protocol produced the pre-master secret in older versions of TLS (pre-TLS 1.2). To derive the final secret in TLS, further nonces were exchanged between the parties (to ensure that both parties were alive and the key was fresh). Then, a master secret was derived from the pre-master secret and the nonces. Finally, key confirmation was provided by the entire protocol transcript being hashed and encrypted under the master secret (the so-called *finished* message). In TLS, the resulting key is not indistinguishable from random as the encrypted finished message provides the adversary with a trivial check to determine whether a key is real or not. However, the protocol can be shown to be secure for the purposes of using the master secret to produce a secure bi-directional channel between the server and the client.

A more basic issue with the above protocol is that it is not forward-secure. Any adversary who records a session now, and in the future manages to obtain the server's long-term secret sk , can obtain the pre-master secret, and hence decrypt the entire session.

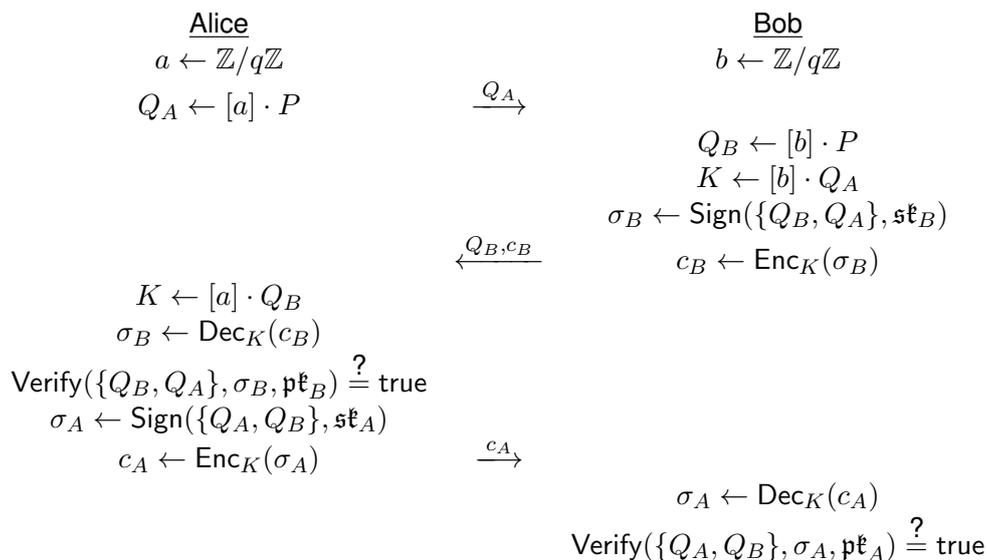
Diffie–Hellman Key Agreement: To avoid the issues with forward secrecy of RSA-based key transport, modern protocols make use of Diffie–Hellman key exchange. This allows two parties to agree on a uniformly random key, which is indistinguishable from random assuming the Decision Diffie–Hellman problem is hard

$$\begin{array}{ccc}
 \text{Alice} & & \text{Bob} \\
 a \leftarrow \mathbb{Z}/q\mathbb{Z} & & b \leftarrow \mathbb{Z}/q\mathbb{Z} \\
 Q_A \leftarrow [a] \cdot P & \xrightarrow{Q_A} & \\
 & \xleftarrow{Q_B} & Q_B \leftarrow [b] \cdot P \\
 K \leftarrow [a] \cdot Q_B & & K \leftarrow [b] \cdot Q_A
 \end{array}$$

This protocol provides forward secrecy, but provides no form of authentication. Due to this, the protocol suffers from an obvious man-in-the-middle attack. To obtain mutual authentication, the message flow of Q_A is signed by Alice's public key and the message flow of Q_B is signed by Bob's public key. This prevents the man-in-the-middle attack. However, since the signatures are not bound into the message, the signed-Diffie–Hellman protocol suffers from an unknown-key-share attack; an adversary (Charlie) can strip Alice's signature from Q_A and replace it with their signature. The adversary does not learn the secret, but does convince Bob he is talking to another entity.

The one-way authenticated version of Diffie–Hellman key agreement is the preferred method of key agreement in modern TLS deployments, and is the only method of key agreement supported by TLS 1.3. In TLS, the finished message, which hashes the entire transcript, prevents the above unknown-key-share attack. However, it also prevents the protocol from producing keys which are indistinguishable from random, as mentioned above.

Station-to-Station Protocol: The Station-to-Station (STS) protocol can be used to prevent unknown-key-share attacks on signed Diffie–Hellman and maintain key indistinguishability. In this protocol, the Diffie–Hellman derived key is used to encrypt the signatures, thus ensuring the signatures cannot be stripped off the messages. In addition, the signatures are applied to the transcript so as to convince both receiving parties that the other party is 'alive'.



9 Advanced Protocols

[2, c20–c22]

Modern cryptography has examined a number of more complex protocols to achieve more complex ends. For example, secure e-voting schemes, secure auctions, data storage and retrieval, etc. Most of these advanced protocols are either based on the simpler components described in this section and/or on the encryption and signature schemes with special properties discussed in the next section. Here we restrict ourselves to discussing three standard protocols: Oblivious Transfer, Zero-Knowledge and Multi-Party Computation.

9.1 Oblivious Transfer

While Oblivious Transfer is at the heart of many advanced protocols, it is a surprisingly simple primitive which enables one to accomplish various more complex tasks. In the following, we describe the basic 1-out-of-2 Oblivious Transfer, but extensions to n -out-of- m are immediate. In all cases, the protocol is one executed between a Sender and a Receiver.

In a 1-out-of-2 Oblivious Transfer, the Sender has two messages m_0 and m_1 , whilst the Receiver has an input bit b . The output of the protocol should be the message m_b for the Receiver, and the Sender obtains no output. In particular, the Receiver learns nothing about the message m_{1-b} , whilst the Sender learns nothing about the bit b .

This passively secure protocol can be implemented as follows. We assume the Sender's messages are two elements M_0 and M_1 in an elliptic curve group $E(\mathbb{F}_p)$ of prime order q .

<u>Sender</u>		<u>Receiver</u>
$C \leftarrow E(\mathbb{F}_p)$	\xrightarrow{C}	
		$x \leftarrow (\mathbb{Z}/q\mathbb{Z})$
		$Q_b \leftarrow [x] \cdot P$
		$Q_{1-b} \leftarrow C - Q_b$
		$\xleftarrow{Q_0}$
$Q_1 \leftarrow C - Q_0$		
$k \leftarrow (\mathbb{Z}/\mathbb{Z})$		
$C_1 \leftarrow [k] \cdot P$		
$E_0 \leftarrow M_0 + [k] \cdot Q_0$		
$E_1 \leftarrow M_1 + [k] \cdot Q_1$		
	$\xrightarrow{C_1, E_0, E_1}$	
		$M_B \leftarrow E_b - [x] \cdot C_1$

The extension to an actively secure protocol is only a little more complex, but beyond the scope of this article.

9.2 Zero-Knowledge

A Zero-Knowledge protocol is a protocol executed between a Prover and a Verifier in which the Prover demonstrates that a statement is true, without revealing why the statement is true. More formally, consider an NP language \mathcal{L} . An interactive proof system for \mathcal{L} is a sequence of protocol executions by an (infinitely powerful) Prover P and a (probabilistic polynomial time) Verifier V , which on joint input x proceeds as follows:

<u>Verifier</u>		<u>Prover</u>
	$\xleftarrow{p_1}$	$(p_1, s'_1) \leftarrow P_1(x)$
$(v_1, s_1) \leftarrow V_1(x, p_1)$	$\xrightarrow{v_1}$	
	$\xleftarrow{p_2}$	$(p_2, s'_2) \leftarrow P_2(s'_1, v_1)$
$(v_2, s_2) \leftarrow V_2(s_1, p_2)$	$\xrightarrow{v_2}$	
	$\xleftarrow{p_3}$	$(p_3, s'_3) \leftarrow P_3(s'_2, v_2)$
	\vdots	\vdots
	$\xrightarrow{p_r}$	$(p_r, s'_r) \leftarrow P_r(s'_{r-1}, v_{r-1})$

By the end of the protocol, the Verifier will output either true or false. An interactive proof system is one which is both *complete* and *sound*

- **Completeness:** If the statement x is true, i.e. $x \in \mathcal{L}$, then if the Prover is honest then the Verifier will output true.

- **Soundness:** If the statement is false, i.e. $x \notin \mathcal{L}$, then no cheating Prover can convince the Verifier with probability greater than p .

Note that even if p is large (say $p = 0.5$) then repeating the proof multiple times can reduce the soundness probability to anything desired. Of course, protocols with small p to start with are going to be more efficient.

For any NP statement, there is a trivial interactive proof system. Namely, the Prover simply sends over the witness w which the Verifier then verifies. However, this reveals the witness. In a zero-knowledge proof, we obtain the same goal, but the Verifier learns nothing bar the fact that $x \in \mathcal{L}$. To formally define zero-knowledge, we insist that there is a (probabilistic polynomial time) *simulator* S which can produce protocol transcripts identical to the transcripts produced between a Verifier and an honest Prover; except the simulator has no access to the Prover. This implies that the Verifier cannot use the transcript to perform any other task, since what it learned from the transcript it could have produced without the Prover by simply running the simulator.

A zero-knowledge proof is said to be *perfect zero-knowledge* if the distribution of transcripts produced by the simulator is identical to those produced between a valid prover and verifier. If the two distributions only cannot be distinguished by an efficient algorithm we say we have *computational zero-knowledge*.

A zero-knowledge proof is said to be a *proof of knowledge* if a Verifier given rewinding access to the prover (i.e. the Verifier can keep resetting the Prover to a previous protocol state and continue executing) can extract the underlying witness w . This implies that the Prover must 'know' w since we can extract w from it.

A *non-interactive zero-knowledge proof* is one in which there is no message flowing from the Verifier to the Prover, and only one message flowing from the Prover to the Verifier. Such non-interactive proofs require additional setup assumptions, such as a Common Reference String (CRS), or they require one to assume the Random Oracle Model.

Σ -Protocols The earlier Σ -protocol for identification is a zero-knowledge proof of knowledge.

$$\begin{array}{ccc}
 \text{Verifier} & & \text{Prover} \\
 & & k \leftarrow \mathbb{Z}/q\mathbb{Z} \\
 & \xleftarrow{r} & R \leftarrow [k] \cdot P \\
 e \leftarrow \mathbb{Z}/q\mathbb{Z} & \xrightarrow{e} & \\
 & \xrightarrow{s} & s \leftarrow (k + e \cdot x) \pmod{q} \\
 R \stackrel{?}{=} [s] \cdot P - e \cdot Q & &
 \end{array}$$

The protocol is obviously complete, and the soundness error is $1/q$. That it is zero-knowledge follows from the following simulation, which first samples $e, s \leftarrow \mathbb{Z}/q\mathbb{Z}$ and then computes $R = [s]P - e \cdot Q$; the resulting simulated transcript being (R, e, s) . Namely, the simulator computes things in the wrong order.

The protocol is also a proof of knowledge since if we execute two protocol runs with the same R value but different e -values (e_1 and e_2) then we obtain two s -values (s_1 and s_2). This is done by rewinding the prover to just after it has sent its first message. If the two obtained transcripts (R, e_1, s_1) and (R, e_2, s_2) are both valid then we have

$$R = [s_1] \cdot P - e_1 \cdot Q = [s_2] \cdot P - e_2 \cdot Q$$

and so

$$[s_1 - s_2] \cdot P = [e_1 - e_2] \cdot Q$$

and hence

$$Q = \begin{bmatrix} s_1 - s_2 \\ e_1 - e_2 \end{bmatrix} \cdot P$$

and hence we 'extract' the secret x from $x = (s_1 - s_2)/(e_1 - e_2) \pmod{q}$.

9.3 Secure Multi-Party Computation

Multi-Party Computation (MPC) is a technique to enable a set of parties to compute on data, without learning anything about the data. Consider n parties P_1, \dots, P_n each with input x_1, \dots, x_n . MPC allows these parties to compute *any* function $f(x_1, \dots, x_n)$ of these inputs without revealing any information about the x_i to each other, bar what can be deduced from the output of the function f .

In an MPC protocol, we assume that a certain set of the parties is corrupt we call this set A . In *statically* secure protocols, this set is defined at the start of the protocol, but remains unknown to the honest parties. In an *adaptively* secure protocol, the set can be chosen by the adversary as the protocol progresses. An MPC protocol is said to be *passively* secure if the parties in A follow the protocol, but try to learn data about the honest parties' inputs from their joint view. In an *actively* secure protocol, the parties in A can arbitrarily deviate from the protocol.

An MPC protocol should be *correct*, i.e. it outputs the correct answer if all parties follow the protocol. It should also be *secure*, i.e. the dishonest parties should learn nothing about the inputs of the honest parties. In the case of active adversaries, a protocol is said to be *robust* if the honest parties will obtain the correct output, even when the dishonest parties deviate from the protocol. A protocol which is not robust, but which aborts with overwhelming probability when a dishonest party deviates, is said to be an actively secure MPC protocol *with abort*.

MPC protocols are categorized by whether they utilize information-theoretic primitives (namely secret sharing), or they utilize computationally secure primitives (such as symmetric-key and public-key encryption). They are also further characterized by the properties of the set A . Of particular interest is when the size t of A is bounded by a function of n (so-called threshold schemes). The cases of particular interest are $t < n$, $t < n/2$, and $t < n/3$; the threshold cases of $t < n/2$ and $t < n/3$ can be generalized to Q_2 and Q_3 access structures, as discussed in Section 3.2.

In the information-theoretic setting, one can achieve passively secure MPC in the case of $t < n/2$ (or Q_2 access structures). Actively secure robust MPC is possible in the information-theoretic setting when we have $t < n/3$ (or Q_3 access structures). All of these protocols are achieved using secret sharing schemes.

In the computational setting, one can achieve actively secure robust computation when $t < n/2$, using Oblivious Transfer as the basic computational foundation. The interesting case of two party computation is done using the Yao protocol. This protocol has one party (the Circuit Creator, also called the Garbler) 'encrypting' a boolean function gate by gate using a cipher such as AES, the circuit is then sent to the other party (called the Circuit Evaluator). The Evaluator then obtains the 'keys' for their input values from the Creator using Oblivious Transfer, and can then evaluate the circuit.

Modern MPC protocols have looked at active security with abort in the case of $t < n$. The modern protocols are divided into a function-dependent offline phase, which requires public key functionality but which is function independent, then a function-dependent online phase which mainly uses information-theoretic primitives. Since information theoretic primitives are usually very fast, this means the time-critical online phase can be executed as fast as possible.

10 Public Key Encryption/Signatures With Special Properties

[1, c13]

A major part of modern cryptography over the last twenty years has been the construction of encryption and signature algorithms with special properties or advanced functionalities. A number of the

following have been deployed in specialized systems (for example, U-PROVE, IDEMIX, attestation protocols and some crypto-currencies). We recap the main variants below, giving for each one the basic idea behind their construction.

Group Signatures A group signature scheme defined a group public key pk , associated to a number of secret keys sk_1, \dots, sk_n . The public key is usually determined by an entity called a *Group Manager*, during an interaction with the group members. Given a group signature s , one cannot tell which secret key signed it, although one is guaranteed that one did. Thus group signatures provide the anonymity of a Signer. Most group signature algorithms have a special entity called an *Opener* who has some secret information which enables them to revoke the anonymity of a Signer. This last property ensures one can identify group members who act dishonestly in some way.

A group signature scheme can either support *static* or *dynamic* groups. In a static group signature scheme, the group members are fixed at the start of the protocol, when the public key is fixed. In a dynamic group signature scheme the group manager can add members into the group as the protocol proceeds, and (often) revoke members as well.

An example of this type of signature scheme which is currently deployed is the Direct Anonymous Attestation (DAA) protocol; which is essentially a group signature scheme in which the *Opener* is replaced with a form of user controlled linkability; i.e. a signer can decide whether two signatures output by the specific signer can be linked or not.

Ring Signatures A ring signature scheme is much like a group signature scheme, but in a ring signature there is no group manager. Each user in a ring signature scheme has a public/private key pair (pk_i, sk_i) . At the point of signing, the Signer selects a subset of the public keys (containing this own), which is called a ring of public keys, and then produces a signature. The Receiver knows the signature was produced by someone in the ring, but not which member of the ring.

Blind Signatures A blind signature scheme is a two party protocol in which a one party (the User) wants to obtain the signature on a message by a second party (the Signer). However, the Signer is not allowed to know which message is being signed. For example, the Signer may be simply notarising that something happened, but does not need to know precisely what. Security requires that the Signer should not learn anything about any message passed to it for signing, and the user should not obtain the signature on any message other than those they submitted for signing.

Identity-Based Encryption In normal public key encryption, a user obtains a public key pk , along with a certificate C . The certificate is produced by a trusted third party, and binds the public key to the identity. Usually, a certificate is a digitally signed statement containing the public key and the associated user identity. So, when sending a message to Alice the Sender is sure that Alice is the legitimate holder of public key pk .

Identity Based Encryption (IBE) is an encryption scheme which dispenses with the need for certificate authorities, and certificates. To encrypt to a user, say Alice, we simply use her identity *Alice* as the public key, plus a global 'system' public key. However, to enable Alice to decrypt, we must have a trusted third party, called a Key Generation Centre, which can provide Alice with her secret key. This third party uses its knowledge of the 'system' secret key to be able to derive Alice's secret key. Whilst dispensing with certificates, an IBE system inherently has a notion of key escrow; the Key Generation Centre can decrypt all messages.

Linearly Homomorphic Encryption In a linearly homomorphic encryption scheme one can perform a number of linear operations on ciphertexts, which result in a ciphertext encrypting a message having had the same operations performed on the plaintext. Thus, given two encryptions $c_1 \leftarrow \text{Enc}(m_1, pk; r_1)$ and $c_2 \leftarrow \text{Enc}(m_2, pk; r_2)$ one can form a 'sum' operation $c \leftarrow c_1 \oplus c_2$ such that c decrypts to $m_1 + m_2$. The standard example of such encryption schemes is the Paillier en-

encryption scheme, which encrypts elements $m \in (\mathbb{Z}/N\mathbb{Z})$, for an RSA-modulus N by computing $c \leftarrow (1 + N)^m \cdot r^N \pmod{N^2}$ where r is selected in $\mathbb{Z}/N\mathbb{Z}$.

Such encryption algorithms can never be IND-CCA secure, as the homomorphic property produces a trivial malleability which can be exploited by a CCA attacker. However, they can have applications in many interesting areas. For example, one can use a linearly homomorphic encryption scheme to add up votes in a digitally balloted election for two candidates, where each vote is an encryption of either the message zero or one.

Fully Homomorphic Encryption Fully Homomorphic Encryption (or FHE) is an extension to linearly homomorphic encryption, in that one can not only homomorphically evaluate linear functions, but also non-linear ones. In particular, the ability to homomorphically evaluate both addition *and* multiplication on encrypted data enables one to (theoretically) evaluate any function. Applications of FHE which have been envisioned are things such as performing complex search queries on encrypted medical data etc. Thus, FHE is very interesting in a cloud environment.

However, all existing FHE schemes are highly inefficient. Resulting in only very simple functions being able to be evaluated in suitable time limits. A scheme which can perform homomorphic operations from a restricted class of functions (for example, to homomorphically evaluate all multi-variate polynomials of total degree five) is called a Somewhat Homomorphic Encryption (or SHE) scheme. Obviously, if the set of functions are all multi-variate polynomials of degree one, then the SHE scheme is a linear homomorphic encryption scheme.

CROSS-REFERENCE OF TOPICS VS REFERENCE MATERIAL

	KatzLindell [1]	Smart [2]
1 Mathematics	c8–c9, App B	c1–c5
2 Cryptographic Security Models	c1–c4	c11
4 Symmetric Primitives	c3–c6	c11–c14
3 Information-theoretically Secure Constructions	c2	c19
5 Symmetric Encryption and Authentication	c3–c4	c13–c14
6 Public Key Encryption	c11	c15–c17
7 Public Key Signatures	c12	c16
8 Standard Protocols	–	c18
9 Advanced Protocols	–	c20–c22
10 Public Key Encryption/Signatures With Special Properties	c13	–

FURTHER READING

The following two text books are recommended to obtain further information on the topics in this knowledge area.

Introduction to Modern Cryptography (J. Katz and Y. Lindell) [1]

A standard modern textbook covering aspects of the design of cryptographic schemes from a provable security perspective.

Cryptography Made Simple (N.P. Smart) [2]

A textbook with less mathematical rigour than the previously mentioned one, but which also covers a wider range of areas (including zero-knowledge and MPC), and touches on aspects related to implementation.

REFERENCES

- [1] J. Katz and Y. Lindell, *Introduction to Modern Cryptography, Second Edition*. CRC Press, 2014.
[2] N. P. Smart, *Cryptography Made Simple*, ser. Information Security and Cryptography. Springer, 2016. [Online]. Available: <https://doi.org/10.1007/978-3-319-21936-3>

ACRONYMS

AEAD Authenticated Encryption with Authenticated Data. 10, 11

AES Advanced Encryption Standard. 5, 6, 8, 23

CBC Cipher Block Chaining. 10, 11

CCA Chosen Ciphertext Attack. 4, 6, 13–15

CFB Cipher Feedback. 10

CMA Chosen Message Attack. 4

CPA Chosen Plaintext Attack. 4, 6, 13

CRS Common Reference String. 6, 22

CTR Counter Mode. 10, 11

CVP Closest Vector Problem. 5

DAA Direct Anonymous Attestation. 24

DDH Decision Diffie–Hellman. 5, 15

DEM Data Encryption Mechanism. 11, 13

DES Data Encryption Standard. 8

DHP Diffie–Hellman Problem. 5

DLP Discrete Logarithm Problem. 5

DSA Digital Signature Algorithm. 17

ECB Electronic Code Book. 10

ECC Elliptic Curve Cryptography. 15

ECIES Elliptic Curve Integrated Encryption Scheme. 15

FDH Full Domain Hash. 15

FHE Fully Homomorphic Encryption. 25

GCM Galois Counter Mode. 11

GMAC Galois Message Authentication Code. 11

HMAC Hash MAC. 9, 11

IBE Identity Based Encryption. 24

IND Indistinguishable. 4, 6, 13–15

KDF Key Derivation Function. 10, 12

KEM Key Encapsulation Mechanism. 13–15

KMAC Keccak MAC. 12

LWE Learning With Errors. 15

MAC Message Authentication Code. 3–5, 11

MPC Multi-Party Computation. 23

NIST National Institute of Standards and Technology. 15

OAEP Optimized Asymmetric Encryption Padding. 14

OFB Output Feedback. 10

OW One-Way. 3, 4

PASS Passive Attack. 3, 4, 6, 13

PKCS Public Key Cryptography Standards. 16

PKI Public Key Infrastructure. 6

PQC Post-Quantum Cryptography. 5

PRF Pseudo-Random Function. 5

PRP Pseudo-Random Permutation. 5, 6, 8

PSS Probabilistic Signature Scheme. 16

RSA Rivest-Shamir-Adleman. 5, 6, 14, 15

SHE Somewhat Homomorphic Encryption. 25

STS Station-to-Station. 20

SVP Shortest Vector Problem. 5

TLS Transport Layer Security. 19

UC Universal Composability. 6

UF Universal Forgery. 4

XOF Extendable Output Function. 10, 12