# Subgroup membership problems and public key cryptosystems 

Dr. ing. thesis

## Preface

This thesis contains my research work done between January 2000 and June 2004, while employed by the Department of Mathematical Sciences at the Norwegian University of Science and Technology. My advisors have been Idar Hansen and Alexei Rudakov. I am grateful for their continuous help, advice and encouragement.

In January 2004 I visited Ronald Cramer in Aarhus. I would like to thank him and the cryptology group in Aarhus for their hospitality and for the many enlightening discussions. The visit was very helpful for me and for the completion of my work.

Finally, I would like to thank my wife and my daughter for being there for me.

Trondheim, May 2004
Kristian Gjøsteen

## Contents

1 Introduction ..... 1
2 Preliminaries ..... 3
2.1 Notation ..... 3
2.2 Computational model ..... 5
2.3 Common problems ..... 9
2.3.1 Factoring ..... 9
2.3.2 Discrete logarithms ..... 10
2.3.3 The Diffie-Hellman problems ..... 10
2.4 Elementary results ..... 12
3 Provable security ..... 15
3.1 Public key cryptosystems ..... 15
3.2 Adversary models ..... 16
3.3 Security goals ..... 17
3.4 Security notions ..... 20
3.5 Key encapsulation methods ..... 21
3.6 Models for provable security ..... 23
3.6.1 Random oracle model ..... 23
3.6.2 Generic model ..... 24
3.7 Plaintext aware encryption ..... 25
4 Subgroup membership problems ..... 29
4.1 Subgroup membership problems ..... 29
4.2 Splitting problems ..... 32
4.3 Subgroup discrete logarithm problems ..... 33
4.4 Symmetric subgroup membership problems ..... 38
4.5 A catalogue ..... 42
4.5.1 Quadratic Residue problem ..... 42
4.5.2 Higher Residue problem ..... 43
4.5.3 Decision Diffie-Hellman problem ..... 45
4.5.4 Decision Composite Residuosity problem ..... 46
4.5.5 Okamoto-Uchiyama ..... 48
4.5.6 Groups of known order ..... 48
4.5.7 Groups of unknown order ..... 49
4.6 Further reductions ..... 50
5 Homomorphic cryptosystems ..... 53
5.1 The general cryptosystem ..... 53
5.2 Non-standard assumptions ..... 59
5.2.1 Knowledge-of-exponent ..... 59
5.2.2 A security proof ..... 61
5.3 A second homomorphic cryptosystem ..... 62
6 Key encapsulation methods ..... 67
6.1 Security against passive attacks ..... 67
6.2 Security against active attacks ..... 70
7 A secure cryptosystem ..... 75
7.1 Hash proof systems ..... 75
7.2 The general construction ..... 76
7.3 Concrete instances ..... 81
8 Concluding remarks ..... 87

## List of Figures

3.1 Encryption and decryption algorithms for hybrid cryptosystems. ..... 23
4.1 Experiments for the proof of Theorem 4.11. ..... 40
5.1 The cryptosystem $\Pi_{1}$ ..... 54
5.2 The cryptosystem $\Pi_{1^{\prime}}$ ..... 56
5.3 The cryptosystem $\Pi_{2}$. ..... 63
6.1 The key encapsulation method $\Lambda_{1}$ ..... 68
6.2 Experiments for the proof of Theorem 6.1 ..... 68
6.3 The key encapsulation method $\Lambda_{2}$ ..... 71
7.1 The cryptosystem $\Pi_{3}$ ..... 77
7.2 An instantiation of the cryptosystem $\Pi_{3}$ ..... 84

## Chapter 1

## Introduction

Public key encryption was first proposed by Diffie and Hellman [16], and widely popularised with the RSA cryptosystem [37]. Over the years, the security goals of public key encryption have been studied [17, 22], as have adversary models [30, 36], and many public key cryptosystems have been proposed and analysed.

It turns out that the security of many of those cryptosystems $[16,18,22,29$, $34,35]$ are based on a common class of mathematical problems, called subgroup membership problems. Cramer and Shoup [10] designed a chosen-ciphertextsecure cryptosystem based on a general subgroup membership problem (generalising their previous work [9]), and provided two new instances. Yamamura and Saito [41] defined a general subgroup membership problem, catalogued several known subgroup membership problems, and designed a private information retrieval system based on a subgroup membership problem. Nieto, Boyd and Dawson [31] designed a cryptosystem based on essentially a symmetric subgroup membership problem (see Section 4.4 and Section 6.1).

Chapter 2 and 3 contain certain preliminary discussions necessary for the later work.

In Chapter 4, we discuss subgroup membership problems, both abstractly and concrete families. For all of the concrete examples, there is a related problem called the splitting problem. We discuss various elementary reductions, both abstract and for concrete families. In cryptographic applications, a third related problem, called the subgroup discrete logarithm problem, is also interesting, and we discuss this in some detail.

We also discuss a variant of the subgroup membership problem where there are two subgroups that are simultaneously hard to distinguish. We prove a useful reduction (Theorem 4.11) for this case. The technique used in the proof is reused throughout the thesis.

In Chapter 5, we discuss two homomorphic cryptosystems, based on trapdoor splitting problems. This gives us a uniform description of a number of homomorphic cryptosystems, and allows us to apply the theory and results of Chapter 4 to the security of those cryptosystems.

Using the technique of Theorem 4.11, we develop a homomorphic cryptosystem that is not based on a trapdoor problem. This gives us a fairly efficient cryptosystem, with potentially useful properties.

We also discuss the security of a homomorphic cryptosystem under a nonstandard assumption. While these results are very weak, they are stronger than results obtained in the generic model.

In Chapter 6, we develop two key encapsulation methods. The first can be proven secure against passive attacks, using the same technique as in the proof of Theorem 4.11. The second method can be proven secure against active attacks in the random oracle model, but to do this, we need a certain non-standard assumption.

Finally, in Chapter 7 we discuss a small extension to the framework developed by Cramer and Shoup [10], again by essentially reusing the technique used to prove Theorem 4.11. This gives us a cryptosystem that is secure against chosen ciphertext attacks, without recourse to the random oracle model or nonstandard assumptions. The cryptosystem is quite practical, and performs quite well compared to other variants of the Cramer-Shoup cryptosystem.

## Chapter 2

## Preliminaries

We discuss the notation and computational model used throughout this text. We shall also discuss certain computational problems common in cryptography. To simplify the main exposition, we have included a number of elementary results in this chapter.

### 2.1 Notation

The symbols $\mathbb{N}, \mathbb{Z}, \mathbb{Q}$ and $\mathbb{R}$ represent the natural numbers, the integers, the rational numbers and the real numbers, respectively. A $k$-bit integer is an integer $n$ such that $2^{k-1} \leq n<2^{k}$. For an integer $n>1$, we use the notation $x \bmod n$ to denote the remainder of $x$ divided by $n$. We denote the ring of integers modulo $n$ by $\mathbb{Z}_{n}$, and the field with $p$ elements by $\mathbb{F}_{p}$. We use the notation $[a]$ and $[a]_{n}$ to denote the residue class in $\mathbb{Z}_{n}$ represented by the integer $a$. The map $a \mapsto[a]_{n}$ is a ring homomorphism from $\mathbb{Z}$ onto $\mathbb{Z}_{n}$.

A smooth integer is an integer whose prime factors are small. Sometimes, we say $B$-smooth to mean that all the prime factors are below the integer bound $B$, or $S$-smooth to mean that all the prime factors are in the set $S$.

Let $n>0$ be an integer. The Euler phi-function $\phi(n)$ is the number of non-negative integers less than $n$ that are relatively prime to $n$.

If $G$ is a group and $g \in G,\langle g\rangle$ denotes the cyclic subgroup of $G$ generated by $g$.

We say that a function $f: \mathbb{N} \rightarrow \mathbb{R}$ is negligible if for any polynomial function $p(\tau)$, there exists a constant $\tau_{0} \geq 0$ such that

$$
|f(\tau)| \leq \frac{1}{p(\tau)} \text { for } \tau>\tau_{0}
$$

Any function $f$ that is not negligible is significant. We also say that $f(\tau)$ is negligible in $\tau$.

A distribution on a countable set $S$ is a function $\mu: S \rightarrow[0,1]$ such that $\sum_{s \in S} \mu(s)=1$. A pair $(S, \mu)$ is called a probability space. Let $X=(S, \mu)$ be a
probability space. For a subset $S^{\prime}$ of $S$ use the notation

$$
\operatorname{Pr}\left[X \in S^{\prime}\right]=\sum_{s \in S^{\prime}} \mu(s) .
$$

We also use the notation $\operatorname{Pr}[X=s]=\operatorname{Pr}[X \in\{s\}]$. If $X=(S, \mu)$ is a probability space, we shall say that $X$ is a probability space on $S$. When we say that $X$ is distributed according to some distribution $\mu$ on $S$, we mean that $X$ is the probability space $(S, \mu)$.

If $X=(S, \mu)$ is a probability space, then $x \leftarrow X$ means that $x$ is sampled from $S$ according to the distribution $\mu$, or $\operatorname{Pr}[x=s \mid x \leftarrow X]=\operatorname{Pr}[X=s]=\mu(s)$ for all $s \in S$. If $S$ is a finite set, the notation $x \leftarrow S$ means that $x$ is sampled uniformly at random from $S$.

If $S, S^{\prime}$ are countable sets and $f: S \rightarrow S^{\prime}$ is a function, then $f$ induces a map $f^{*}$ from the set of distributions on $S$ to the set of distributions on $S^{\prime}$. If $\mu$ is a distribution on $S$, then the distribution $f^{*}(\mu)=\mu^{\prime}$ on $S^{\prime}$ is given by

$$
\mu^{\prime}\left(s^{\prime}\right)=\sum_{s \in f^{-1}\left(s^{\prime}\right)} \mu(s)
$$

If $X=(S, \mu)$ is a probability space and $f: S \rightarrow S^{\prime}$ is a function, then we denote the probability space $\left(S^{\prime}, \mu^{\prime}\right)$ by $f(X)$.

A triple $\left(S, \mu, f: S \rightarrow S^{\prime}\right)$, where $(S, \mu)$ is a probability space and $f$ is a function, is called a random variable on the probability space $(S, \mu)$ with values in $S^{\prime}$. If we have random variables $X_{1}, X_{2}, \ldots, X_{l}$ and a function $f: S_{1}^{\prime} \times S_{2}^{\prime} \times$ $\cdots \times S_{l}^{\prime} \rightarrow S^{\prime \prime}$, we get a random variable $f\left(X_{1}, X_{2}, \ldots, X_{l}\right)$ on the joint probability space $\left(S_{1}, \mu_{1}\right) \times \cdots \times\left(S_{l}, \mu_{l}\right)$ with values in $S^{\prime \prime}$.

In general, we shall identify the random variable with the resulting probability space.

Let $X$ and $Y$ be probability spaces on some finite set $S$. The statistical distance between $X$ and $Y$ is defined to be

$$
\operatorname{Dist}(X, Y)=\frac{1}{2} \sum_{s \in S}|\operatorname{Pr}[X=s]-\operatorname{Pr}[Y=s]| .
$$

We note that $\operatorname{Dist}(\cdot, \cdot)$ is a metric on the set of probability spaces on $S$, and that $0 \leq \operatorname{Dist}(X, Y) \leq 1$. We say that two probability spaces $X$ and $Y$ are $\epsilon$-close if $\operatorname{Dist}(X, Y) \leq \epsilon$. If two probability spaces $X$ and $Y$ are $\epsilon$-close to each other for some small $\epsilon$, we say that $X$ is almost equal to $Y$, and vice versa. Two families of probability spaces $X(\tau)$ and $Y(\tau)$ are statistically indistinguishable if $\operatorname{Dist}(X(\tau), Y(\tau))$ is negligible in $\tau$.

If $X$ is a probability space on some finite subset $S$ of the real numbers, then the expected value of $X$ is

$$
\mathrm{E}[X]=\sum_{s \in S} \operatorname{Pr}[X=s] s
$$

The projective plane over a field $\mathbb{F}$ is the set of triples $(x: y: z)$, where $x, y$ and $z$ are in the algebraic closure of $\mathbb{F}$, and at least one of $x, y$ and $z$ are nonzero, subject to the equivalence relation that $(x, y, z) \sim\left(x^{\prime}, y^{\prime}, z^{\prime}\right)$ if there exists $\lambda \in \mathbb{F}^{*}$ such that $(x, y, z)=\left(\lambda x^{\prime}, \lambda y^{\prime}, \lambda z^{\prime}\right)$. We denote the equivalence classes by ( $x: y: z$ ) and call them points.

An elliptic curve [40] defined over a field $\mathbb{F}$ of characteristic different from 2 and 3 is the set of points in the projective plane satisfying

$$
y^{2} z=x^{3}+a x z^{2}+b z^{3}
$$

where $a, b \in \mathbb{F}$ and $4 a^{3}-27 b^{2} \neq 0$. The point $(0: 1: 0)$ is called the point at infinity, and it is the only point with $z$-coordinate equal to 0 . The elliptic curve is a group under a geometrically defined operation, where $(0: 1: 0)$ is the identity. The set of $\mathbb{F}$-rational points (points that can be represented by a triple $(x: y: z)$, with $x, y, z \in \mathbb{F})$ is a subgroup.

If $\mathbb{F}$ is a finite field with $p^{k}$ elements, it can be shown that the map $(x: y: z) \mapsto$ $\left(x^{p^{k}}: y^{p^{k}}: z^{p^{k}}\right)$ is an endomorphism. It is called the Frobenius endomorphism.

If $n$ is a composite number, we can define an elliptic curve over the ring $\mathbb{Z}_{n}$ to be the set of points

$$
E\left(\mathbb{Z}_{n}\right)=\left\{(x, y, z) \in \mathbb{Z}^{3} \backslash n \mathbb{Z}^{3} \mid y^{2} z \equiv x^{3}+a x z^{2}+b z^{3} \quad(\bmod n)\right\}
$$

subject to the equivalence relation that $(x, y, z) \sim\left(x^{\prime}, y^{\prime}, z^{\prime}\right)$ if there exists $\lambda \in \mathbb{Z}$ such that $\operatorname{gcd}(\lambda, n)=1$ and

$$
(x, y, z) \equiv\left(\lambda x^{\prime}, \lambda y^{\prime}, \lambda z^{\prime}\right) \quad(\bmod n),
$$

and where $a, b$ are integers such that $4 a^{3}-27 b^{2} \not \equiv 0(\bmod n)$. We denote the equivalence classes by $(x: y: z)$ and call them points. The special point $(0: 1: 0)$ is on the curve.

It can be shown ([20, 25] and the references therein) that the usual formulae for the elliptic curve group operation over fields, suitably modified for operations in rings, actually gives a group operation on the set $E\left(\mathbb{Z}_{n}\right)$, and the point $(0: 1: 0)$ is the identity for this group operation.

If $n=p q$, we can of course consider $E$ to be an elliptic curve over the fields $\mathbb{F}_{p}$ and $\mathbb{F}_{q}$, and using the Chinese remainder theorem, it is easy to show that there is a group isomorphism such that

$$
E\left(\mathbb{Z}_{n}\right) \simeq E\left(\mathbb{F}_{p}\right) \times E\left(\mathbb{F}_{q}\right) .
$$

### 2.2 Computational model

An algorithm is a Turing machine that halts after a finite number of steps, regardless of input. It has an input tape and an output tape. The algorithm may have access to a random tape, in which case we call the algorithm probabilistic. The random tape is just a tape containing random zeros and ones, and is sometimes
called the coin tosses of the algorithm. We usually consider the random tape to be infinitely long. A probabilistic algorithm should terminate after a finite number of steps, regardless of input and random tape.

Sometimes, we want to fix the random tape, and we denote by $A(x ; r)$ the function computed by $A$ when given the fixed random tape $r$. Running an algorithm twice with the same random tape is sometimes called rewinding the random tape.

For a fixed input $x$, we can assume that the random tape is of finite length. The uniform distribution on the set of fixed-length random tapes is then a probability space, and the algorithm $A$ and the fixed input $x$ gives a function on this probability space. We denote this random variable by $A(x)$. If $X$ is a probability space on a finite set of possible inputs, we denote by $A(X)$ the random variable given by applying the function $A(x ; r)$ to the joint probability space of $X$ and algorithm's coin tosses.

An algorithm may have access to one or more oracles. While the word oracle suggests divine origin, our oracles are devices we use to give our algorithms certain well-defined powers of computation. The algorithm may ask a question, and the oracle answers the question in some well-defined way.

The idea is that the oracle performs some computation that the algorithm cannot do on its own, perhaps because we want to allow an algorithm to perform some computation based upon some knowledge ${ }^{1}$ (such as a secret key), without giving that knowledge to the algorithm.

An oracle is modelled as a pair of tapes. The algorithm writes queries to one tape (the query-tape) and reads answers from the other tape (the answer-tape). From the algorithm's point of view, the answer to a query is available immediately after the query was written. Except for the answer, the oracle's computations are hidden from the algorithm.

Every oracle has access to its own random tape which is independent of all other random tapes in the experiment, and the algorithm is not allowed access to the random tape, neither to look at it or to manipulate it. We also allow oracles to preserve state information between queries (that is, the oracle's answer may depend on previous queries to the oracle).

The cost of an algorithm is the number of steps the algorithm needs to run. If the algorithm has access to any oracles, the cost should include the number of queries to each oracle. If the cost is relatively small, we say that the algorithm is efficient.

We say that a function $f: S_{1} \rightarrow S_{2}$ is easy to compute if some deterministic and efficient algorithm for computing $f$ can easily be derived from the description of $f$. We identify $f$ with this algorithm.

[^0]These definition are rather vague and informal. Asymptotically, there is a more precise definition of relatively small cost. An algorithm is polynomial-time if the Turing machine halts after a number of steps polynomial in the length of its input. This also implies that the number of queries to any oracles is bounded by a polynomial in the length of the input. A probabilistic algorithm that is polynomial-time is called efficient.

A problem $\mathcal{P}=(P, X, S, f)$ consists of a set $P$ of problem instances, a probability space $X$ on that set, a finite set $S$ of possible answers, and a function $f: P \rightarrow S$ that gives the answer to every problem instance. To simplify the exposition, we denote sampling a problem instance according to $X$ by $x \leftarrow \mathcal{P}$. The correct answer to instance $x$ is denoted by $\mathcal{P}(x)$.

A subproblem $\mathcal{P}^{\prime} \subseteq \mathcal{P}$ is a tuple ( $P^{\prime}, X^{\prime}, S^{\prime}, f^{\prime}$ ) where $P^{\prime} \subseteq P, X^{\prime}$ is the probability space $X$ restricted to $P^{\prime}, f^{\prime}$ is $f$ restricted to $P^{\prime}$ and $S^{\prime}=S$.

The success probability of an algorithm $A$ in solving a problem $\mathcal{P}$ is

$$
\operatorname{Succ}_{A}^{\mathcal{P}}=\operatorname{Pr}[A(x)=\mathcal{P}(x) \mid x \leftarrow \mathcal{P}] .
$$

The probability is computed over the joint probability space of the algorithm's random tape (if it is probabilistic) and the input.

Since the set of possible answers to a problem is finite, there is always a probability that an algorithm will answer correctly by guessing. The advantage of an algorithm measures how much better it is at answering the problem than an algorithm that guesses an answer according to some fixed distribution. The definition is

$$
\begin{aligned}
\operatorname{Adv}_{A}^{\mathcal{P}}=\inf _{Y} & \mid \operatorname{Pr}[A(x)=\mathcal{P}(x) \mid x \leftarrow \mathcal{P}]- \\
& \operatorname{Pr}[y=\mathcal{P}(x) \mid x \leftarrow \mathcal{P} \wedge y \leftarrow Y] \mid,
\end{aligned}
$$

where the distribution $Y$ ranges over all possible distributions on the set of possible answers to $\mathcal{P}$. Note that this definition of advantage may not always be appropriate.

A reduction from a problem $\mathcal{P}$ to a problem $\mathcal{P}^{\prime}$ is an algorithm $A$ that uses an oracle for solving $\mathcal{P}^{\prime}$ to solve $\mathcal{P}$. The cost of the reduction is the cost of running $A$. A reduction $A$ is only interesting if its cost is relatively small, and the success probability of $A$ is not small compared to the success probability of the oracle.

A problem $\mathcal{P}=(P, X, S, f)$ is self-reducible if there is an algorithm $A$ (with a small cost) that takes as input an instance $x$ of $\mathcal{P}$ and outputs an instance $x^{\prime}$ of $\mathcal{P}$ along with a function $g: S \rightarrow S$, such that $g\left(f\left(x^{\prime}\right)\right)=f(x)$, and $A(x)$ is, except with small probability, independent of $x$ and $\epsilon$-close to $X$ for some small $\epsilon$.

Any self-reducible problem that is easy to solve on some significant subset (probability-wise) of its problem instances is easy to solve for almost all of its instances.

A parameterised problem $\mathcal{P}$ is a family $\left\{\mathcal{P}_{\tau}\right\}$ of problems indexed by integers $\tau>0$. For an algorithm $A$, we define success probability as

$$
\operatorname{Succ}_{A}^{\mathcal{P}}(\tau)=\operatorname{Pr}\left[A(x)=\mathcal{P}_{\tau}(x) \mid x \leftarrow \mathcal{P}_{\tau}\right],
$$

and advantage as

$$
\begin{align*}
& \operatorname{Adv}_{A}^{\mathcal{P}}(\tau)=\inf _{Y} \mid \operatorname{Pr}\left[A(x)=\mathcal{P}_{\tau}(x) \mid x \leftarrow \mathcal{P}_{\tau}\right]-  \tag{2.1}\\
& \operatorname{Pr}\left[y=\mathcal{P}_{\tau}(x) \mid x \leftarrow \mathcal{P}_{\tau} \wedge y \leftarrow Y\right] \mid,
\end{align*}
$$

where the distribution $Y$ ranges over all possible distributions on the set of possible answers to $\mathcal{P}_{\tau}$.

We drop the word parameterised if it is clear from the context that the problem is parameterised.

Let $\mathcal{P}$ and $\mathcal{P}^{\prime}$ be parameterised problems. A reduction from $\mathcal{P}$ to $\mathcal{P}^{\prime}$ is an algorithm $A$ that reduces every problem in $\mathcal{P}_{\tau}$ to the corresponding problem in $\mathcal{P}_{\tau}^{\prime}$. A reduction is efficient if the algorithm $A$ is efficient in $\tau$, and if the oracle for $\mathcal{P}^{\prime}$ has significant advantage, then $A$ has significant advantage.

If there is an efficient reduction from $\mathcal{P}$ to $\mathcal{P}^{\prime}$, we say that $\mathcal{P}$ reduces to $\mathcal{P}^{\prime}$. If $\mathcal{P}$ reduces to $\mathcal{P}^{\prime}$ and $\mathcal{P}^{\prime}$ reduces to $\mathcal{P}$, we say that $\mathcal{P}$ and $\mathcal{P}^{\prime}$ are equivalent.

In general, we say that a problem $\mathcal{P}$ is $(c, \epsilon)$-hard if there are no algorithms that solve $\mathcal{P}$ with advantage greater than $\epsilon$ at a cost less than $c$. If we simply say hard, we actually mean $(c, \epsilon)$-hard for some large cost $c$ and some small probability $\epsilon$, so that it is extremely unlikely that any adversary has the resources to solve the problem.

We say that a parameterised problem $\mathcal{P}$ is hard if for any efficient algorithm, its advantage is negligible in $\tau$.

Let $\mathcal{P}$ be a parameterised problem and suppose that $A$ is an algorithm that solves $\mathcal{P}$. $A$ may have access to some oracles. Suppose $A^{\prime}$ is another algorithm that solves $\mathcal{P}$, that it has access to the same oracles as $A$ has, but that it not only outputs the answer to $\mathcal{P}$, but also outputs a trace tr that contains some information about its calculations. We shall need an experiment in order to define what a trace-variant is.

## Experiment 1.

Input: $\mathcal{P}, A, A^{\prime}, \tau$.

1. $x \leftarrow \mathcal{P}_{\tau}$.
2. $r \leftarrow\{0,1\}^{*}$.
3. $y \leftarrow A(x ; r)$.
4. $\left(y^{\prime}, t r\right) \leftarrow A^{\prime}(x ; r)$.
5. If $y=y^{\prime}$, then output 1 , otherwise output 0 .

Output: 0 or 1 .

We say that $A^{\prime}$ is a trace-variant of $A$ if

$$
\operatorname{Pr}\left[\operatorname{Exp} \mathbf{1}\left(\mathcal{P}, A, A^{\prime}, \tau\right)=0\right]
$$

is negligible in $\tau$. Note that $A$ and $A^{\prime}$ are run with the same random tape. If $A$ and $A^{\prime}$ have access to any oracles, then the oracles' random tapes are also rewound.

Obviously, there are all kinds of trivial trace-variants of any algorithm.
A trapdoor problem $\mathcal{P}$ is a set of pairs $\left(\mathcal{P}^{\prime}, \sigma\right)$ along with a distribution on that set, where $\mathcal{P}^{\prime}$ is a problem and $\sigma$ is a function that is easy to compute, such that for any instance $x$ of $\mathcal{P}^{\prime}, \sigma(x)=\mathcal{P}^{\prime}(x)$.

The underlying problem is the problem given by first sampling $\left(\mathcal{P}^{\prime}, \sigma\right)$ from $\mathcal{P}$, forgetting $\sigma$, and then sampling a problem instance $x$ from $\mathcal{P}^{\prime}$.

A parameterised trapdoor problem $\mathcal{P}$ is a family $\left\{\mathcal{P}_{\tau}\right\}$ of trapdoor problems indexed by integers $\tau>0$. We say that a trapdoor problem is hard if the underlying parameterised problem is hard.

### 2.3 Common problems

We shall briefly discuss several problems that are common in cryptography and that are generally assumed to be hard. For a general discussion of these problems, see [28].

### 2.3.1 Factoring

Definition. Let $n>1$ be a composite integer. The factoring problem $\mathcal{F A C} \mathcal{T}_{n}$ is to find the prime factorisation of $n$.

It can be shown that anyone who can compute (a multiple) of $\phi(n)$ can factor.
Let $n>0$ be a composite integer with two prime factors $p$ and $q$, such that both $p$ and $q$ are distinct $\tau$-bit integers. Let $N_{\tau}$ be a probability space on the set of all such integers, and let $N=\left\{N_{\tau}\right\}_{\tau \geq 1}$. The distribution on $N_{\tau}$ is the distribution given by some prime sampling algorithm. It is generally believed that $\mathcal{F A C T}_{N}$ is a hard problem for most reasonable prime sampling algorithms.

Definition. Let $n$ be a composite integer with two primes factors $p$ and $q$. Let $e>1$ be an integer that is relatively prime to $\phi(n)$, and let $f: \mathbb{Z}_{n} \rightarrow \mathbb{Z}_{n}$ be the function defined by $f(m)=m^{e}$.

The $R S A$ [37] problem $\mathcal{R S} \mathcal{A}_{n, e}$ is the following: The instances are elements of $\mathbb{Z}_{n}$ chosen uniformly at random, and the answer to instance $x$ is $f^{-1}(x)$.

To fully specify the RSA problem, we should not only say how $n$ is chosen, but also how $e$ is chosen. Typically, this is the smallest prime $e$ above some bound that satisfies $\operatorname{gcd}(e, \phi(n))=1$. When we do not particularly care about $e$, we ignore it.

We derive the RSA problem $\mathcal{R S} \mathcal{A}_{N}$ from $N$. It is obvious that $\mathcal{R S} \mathcal{A}_{N}$ reduces to $\mathcal{F A C} \mathcal{T}_{N}$. It is not known whether the RSA problem is easier than factoring, but the best known method for computing the inverse function is by factoring the modulus.

The RSA trapdoor problem associates with the problem $\mathcal{R S} \mathcal{A}_{n, e}$ the function $x \mapsto x^{d}$, where $d \equiv e^{-1}(\bmod \phi(n))$.

### 2.3.2 Discrete logarithms

Definition. Let $G$ be a finite cyclic group, let $g$ be a generator, and let $y \in G$. The discrete logarithm of $y$ to the base $g$ is an integer $0 \leq a<|G|$ such that $y=g^{a}$. We sometimes write $\log _{g} y \equiv a(\bmod |G|)$.

The discrete logarithm problem $\mathcal{D} \mathcal{L}_{G}$ is the set of triples $(G, g, y)$, where $y$ is sampled uniformly from $G$ and the answer is $\log _{g} y$.

If $|G|$ is a smooth integer, then the Pohlig-Hellman [28] algorithm shows us that the discrete logarithm problem in $|G|$ is not difficult.

It is easy to show that if one can compute discrete logarithms to some base $g$, then one can compute the order of the group. It is also easy to show that if one can compute discrete logarithms to one base, one can compute discrete logarithms to any base.

Suppose the group order $|G|$ is known. If $X$ is distributed uniformly on $\{0, \ldots,|G|-1\}$, then $y g^{X}$ is distributed uniformly on $G$, it is independent of $y$ and $\log _{g} y g^{X} \equiv X+\log _{g} y(\bmod |G|)$. This means that the discrete logarithm problem is self-reducible when the group order is known.

If $\Gamma$ is a probability space on a set of groups the generalised discrete logarithm problem $\mathcal{D} \mathcal{L}_{\Gamma}$ derived from $\Gamma$ is a set of triples $(G, g, y)$ where $G$ is sampled from $\Gamma$ and $g, y$ are sampled uniformly from $G$. If $y$ is in the subgroup of $G$ generated by $g$, then the answer is $\log _{g} y$, otherwise the answer is the special symbol $\perp$.

Let $N=\left\{N_{\tau}\right\}$ be as in Section 2.3.1. To $N_{\tau}$ we associate the groups $\Gamma_{\tau}=$ $\left\{\mathbb{Z}_{n}^{*} \mid n \in N_{\tau}\right\}$ and we get the family $\Gamma=\left\{\Gamma_{\tau}\right\}$ associated to $N$. It can be shown that $\mathcal{F A C} \mathcal{T}_{N}$ reduces to $\mathcal{D} \mathcal{L}_{\Gamma}$.

### 2.3.3 The Diffie-Hellman problems

Definition. Let $G$ be a cyclic group generated by $g$. The Computational DiffieHellman [16] problem $\mathcal{C D H}_{G}$ is the set of quadruples ( $G, g, x, y$ ), where $x$ and $y$ are sampled uniformly from $G$ and the answer is $z \in G$ such that $\log _{g} z \equiv$ $\log _{g} x \log _{g} y(\bmod |G|)$.

Suppose $G$ has prime order. Let $(G, g, x, y)$ be an instance of $\mathcal{C D} \mathcal{H}_{G}$ with answer $z$. Let $U, V$ and $W$ be uniformly and independently distributed on $\{0, \ldots,|G|-1\}$. Then $\left(G, g^{U}, x^{U V}, y^{U W}\right)$ is the uniform distribution on the set of problem instances for $\mathcal{C D} \mathcal{H}_{G}$, it is independent of $(G, g, x, y)$ (except when $x$ or $y$ is 1 ), and the answer is $z^{U V W}$. This means that $\mathcal{C D H}_{G}$ is self-reducible.

It is clear that anyone who can compute discrete logarithms in $G$ can also solve the Computational Diffie-Hellman problem in $G$. As for discrete logarithms, we
can define a generalised Computational Diffie-Hellman problem. If $x$ or $y$ is not in $\langle g\rangle$, then the correct answer is $\perp$.

Maurer [27] showed that, given a group of prime order $p$ and an elliptic curve defined over $\mathbb{F}_{p}$ with a smooth number of $\mathbb{F}_{p}$-rational points, there is a reasonably efficient reduction from the Computational Diffie-Hellman problem to the discrete logarithm problem. As a consequence, the Computational Diffie-Hellman problem is widely believed to be as difficult as computing discrete logarithms.

A safe prime is a prime $p$ such that $(p-1) / 2$ is also prime. Correspondingly, $p$ is a Sophie-Germain-prime if $2 p+1$ is also prime. Let $N=\left\{N_{\tau}\right\}$ be as in Section 2.3.1, except that we restrict the primes involved to be safe primes. Let $Q_{n}$ be the subgroup of quadratic residues of $\mathbb{Z}_{n}^{*}$, let $\Gamma_{\tau}$ be the set of groups $\left\{Q_{n} \mid n \in N_{\tau}\right\}$ with the derived distribution, and let $\Gamma=\left\{\Gamma_{\tau}\right\}$. It can be shown that $\mathcal{F A C} \mathcal{I}_{N}$ reduces to $\mathcal{C D} \mathcal{H}_{\Gamma}$.

In contrast to factoring and discrete logarithms, there seems to be no easy way to decide if an algorithm returns the correct answer to a Computational DiffieHellman problem. This gives rise to a decision problem.

Definition. Let $G$ be a cyclic group generated by $g$. The Decision DiffieHellman [7] problem $\mathcal{D D H}_{G}$ in $G$ is to distinguish the distribution defined by $\left(g, g^{X}, g^{Y}, g^{X Y}\right)$ from $\left(g, g^{X}, g^{Y}, g^{Z}\right)$, where $X, Y$ are independently and uniformly distributed on $\{0, \ldots,|G|-1\}$, and $Z$ is uniformly distributed on $\{0, \ldots$, $|G|-1\} \backslash\{X Y \bmod |G|\}$.

Suppose $G$ has prime order. It is quite clear that if $(g, x, y, z)$ is an instance of $\mathcal{D D H}_{G}$, then with $U, V, W$ uniformly and independently distributed on $\{0, \ldots,|G|-1\},\left(g^{U}, x^{U V}, y^{U W}, z^{U V W}\right)$ is uniformly distributed over the set of problem instances for $\mathcal{D D} \mathcal{H}_{G}$ and it is independent (except when $x$ or $y$ is 1 ), and the answer remains the same. Hence, the $\mathcal{D D}_{G}$ problem is self-reducible.

We can use the following experiment to measure the advantage of an attacker against the Decision Diffie-Hellman problem.

## Experiment 2.

Input: $A, G, g ; G=\langle g\rangle$.

1. $u, v \leftarrow\{0, \ldots,|G|-1\}$.
2. $b \leftarrow\{0,1\}$.
3. If $b=1$, then $w \leftarrow u v$, otherwise $w \leftarrow\{0, \ldots,|G|-1\} \backslash\{u v \bmod |G|\}$.
4. $b^{\prime} \leftarrow A\left(G, g, g^{u}, g^{v}, g^{w}\right)$.
5. If $b=b^{\prime}$, output 1 , otherwise output 0 .

Output: 0 or 1 .
It is easy to see that

$$
\operatorname{Adv}_{A}^{\mathcal{D} \mathcal{D H}_{G}}=|\operatorname{Pr}[\operatorname{Exp} 2(A, G, g)=1]-1 / 2| .
$$

A nice survey on the Decision Diffie-Hellman problem is [7]. One interesting result is that there does exist groups in which the Computational Diffie-Hellman problem is believed to be difficult, but the Decision Diffie-Hellman problem is easy [24]. We shall return to this in Chapter 6.

There is a trapdoor variant of $\mathcal{C D} \mathcal{H}_{G}$. The problem is given by $\left(G, g, g^{a}\right)$, where $a$ is sampled uniformly from $\{0, \ldots,|G|-1\}$, the problem instance is $\left(G, g, g^{a}, x\right)$, where $x$ is sampled uniformly from $G$, and $\sigma$ is the function $x \mapsto x^{a}$.

### 2.4 Elementary results

We will need several elementary results later on, and to streamline the presentation, we list them here.

Lemma 2.1. Let $S$ be a finite set and let $T \subseteq S$ be a subset. Let $X$ and $Y$ be two probability spaces on $S$ that are $\epsilon$-close. Then

$$
|\operatorname{Pr}[x \in T \mid x \leftarrow X]-\operatorname{Pr}[y \in T \mid y \leftarrow Y]| \leq \epsilon
$$

Proof. Computation.
Lemma 2.2. Let $S^{\prime}$ be a proper, non-empty subset of the finite set $S$. Let $X$ be uniformly distributed on $S$, and let $Y$ be uniformly distributed on $S \backslash S^{\prime}$. Then

$$
\operatorname{Dist}(X, Y)=\frac{2\left|S^{\prime}\right|}{|S|}
$$

Proof. Computation.
Lemma 2.3. Let $n$ be a positive integer, and let $t_{0}$ be an integer such $n<$ $2^{t_{0}}$. Let $X$ be uniformly distributed on $\mathbb{Z}_{n}$, let $Y_{t}$ be uniformly distributed on $\left\{0, \ldots, 2^{t_{0}+t}-1\right\} \subseteq \mathbb{Z}$ and let $Z_{t}$ be the probability space $\left[Y_{t}\right]_{n}$. Then

$$
\operatorname{Dist}\left(X, Z_{t}\right) \leq 2^{-t}
$$

Proof. From the definition, we have that

$$
\operatorname{Dist}\left(X, Z_{t}\right)=\sum_{i=0}^{n-1}\left|\operatorname{Pr}[X=[i]]-\operatorname{Pr}\left[Z_{t}=[i]\right]\right|
$$

Let $r=2^{t_{0}+t} / n$ and $m=2^{t_{0}+t} \bmod n$. Reordering the sum and computing the probabilities, we find that

$$
\operatorname{Dist}\left(X, Z_{t}\right)=\frac{m}{n}\left|1-r^{-1}\lceil r\rceil\right|+\frac{n-m}{n}\left|1-r^{-1}\lfloor r\rfloor\right| .
$$

Since $r-1<\lfloor r\rfloor \leq r \leq\lceil r\rceil<r+1$, we find that

$$
\operatorname{Dist}\left(X, Z_{t}\right) \leq \frac{m}{n} \frac{1}{r}+\frac{n-m}{n} \frac{1}{r}=\frac{n}{2^{t_{0}+t}} \leq 2^{-t}
$$

which concludes the proof.

Lemma 2.4. Let $G$ be a finite cyclic group with generator $g$, and suppose $t_{0}$ is an integer such that $|G|<2^{t_{0}}$. Then for any integer $t>0$ there is an algorithm that samples elements of $G$ from a probability space that is $2^{-t}$-close to the uniform distribution, using $O(t)$ exponentiations in the group.
Proof. The algorithm simply samples $y$ uniformly at random from the set $\{0, \ldots$, $\left.2^{t_{0}+t}-1\right\}$ and computes $g^{y}$. It is clear that this requires the same amount of work as $O(t)$ exponentiations in the group, and the resulting distribution is $2^{-t}$-close to uniform by Lemma 2.3.

Lemma 2.5. Let $G$ be a finite cyclic group. The probability that an element $x$ sampled uniformly at random from $G$ is a generator for $G$ is

$$
\operatorname{Pr}[x \text { is a generator } \mid x \leftarrow G]=\frac{\phi(|G|)}{|G|} .
$$

Proof. $G$ is isomorphic to $\mathbb{Z}_{|G|}$, and the claim follows from the definition of the Euler phi-function.

Suppose we have an experiment that has a fixed number of possible outcomes, numbered from 0 to $d-1$, that we can repeat the experiment such that two different outcomes are independent, and that outcome number $i$ occurs with a fixed probability $\alpha_{i}$. Let $X_{i}$ count the number of times the outcome $i$ occurs in a collection of experiments. Then $X_{i}$ is a random variable, and ( $X_{0}, X_{1}, \ldots, X_{d-1}$ ) is multinomially distributed with parameters $\alpha_{0}, \ldots, \alpha_{d-1}$.
Lemma 2.6. Let $\left(X_{0}, X_{1}, \ldots, X_{d-1}\right)$ be a multinomially distributed random variable with parameters $\alpha_{0}, \alpha_{1}, \ldots, \alpha_{d-1}$. Suppose that $\alpha_{0} \geq \alpha_{i}+\epsilon$ for $i=1, \ldots, d-1$ and $\epsilon>0$. Then after $\left(d^{3}\left(1-\alpha_{0}\right)\right) /\left(\epsilon^{3} \alpha_{0}^{2}(1-\beta)\right)$ experiments,

$$
\operatorname{Pr}\left[X_{0}=\max \left\{X_{0}, \ldots, X_{d-1}\right\}\right] \geq \beta
$$

Proof. Let $l$ be the number of experiments, let $T_{0}$ be the event that $\left|X_{0} / l-\alpha_{0}\right| \leq$ $(d-1) \epsilon / d$, and let $T_{i}$ be the event that $\left|X_{i} / l-\alpha_{i}\right| \leq \epsilon / d, 1 \leq i<d$.

We have that

$$
\operatorname{Pr}\left[X_{0}=\max \left\{X_{i}\right\}\right] \geq \operatorname{Pr}\left[T_{0} \wedge T_{1} \wedge \cdots \wedge T_{d-1}\right] .
$$

If we assume $T_{1} \wedge \cdots \wedge T_{d-1}$, then

$$
\begin{aligned}
\left|X_{0}-l \alpha_{0}\right| & =\left|l-X_{1}-X_{2}-\cdots-X_{d-1}-l\left(1-\alpha_{1}-\cdots-\alpha_{d-1}\right)\right| \\
& \leq\left|X_{1}-l \alpha_{1}\right|+\cdots+\left|X_{d-1}-l \alpha_{d-1}\right| \\
& \leq \frac{d-1}{d} l \epsilon .
\end{aligned}
$$

So $\operatorname{Pr}\left[T_{0} \mid T_{1} \wedge \cdots \wedge T_{d-1}\right]=1$.
Now we consider $\operatorname{Pr}\left[T_{i} \mid T_{i+1} \wedge \cdots \wedge T_{d-1}\right]$, and the marginal distribution of $X_{i}$. It has parameter

$$
\alpha_{i}^{\prime}=\frac{\alpha_{i}}{1-\sum_{j=i+1}^{d-1} \alpha_{j}},
$$

so with $Y_{i}=l-\sum_{j=i+1}^{d-1} X_{j}$ experiments, we get

$$
\begin{aligned}
\operatorname{Pr}\left[T_{i} \mid T_{i+1} \wedge \ldots T_{d-1}\right] & =1-\frac{\alpha_{i}^{\prime}\left(1-\alpha_{i}^{\prime}\right)}{Y_{i} \epsilon^{2} / d^{2}} \\
& =1-\frac{d^{2} \alpha_{i}\left(1-\sum_{j=i}^{d-1} \alpha_{j}\right)}{\epsilon^{2} Y_{i}\left(1-\sum_{j=i+1}^{d-1} \alpha_{j}\right)^{2}} \\
& \geq 1-\frac{d^{2} \alpha_{i}\left(1-\sum_{j=i}^{d-1} \alpha_{j}\right)}{\epsilon^{2} Y_{i} \alpha_{0}^{2}}
\end{aligned}
$$

Under the condition $T_{i+1} \wedge \cdots \wedge T_{d-1}$, we get that

$$
l-\sum_{j=k+1}^{d-1}\left(l \alpha_{j}+l \epsilon / d\right) \leq Y_{i} \leq l-\sum_{j=k+1}^{d-1}\left(l \alpha_{j}-l \epsilon / d\right) .
$$

This means that

$$
\begin{aligned}
\operatorname{Pr}\left[T_{i} \mid T_{i+1} \wedge \ldots T_{d-1}\right] & \geq 1-\frac{d^{2} \alpha_{i}\left(1-\sum_{j=1}^{d-1} \alpha_{j}\right)}{\epsilon^{2} \alpha_{0}^{2} l\left(1-\sum_{j=k+1}^{d-1}\left(\alpha_{j}+\epsilon / d\right)\right)} \\
& \geq 1-\frac{d^{2} \alpha_{i}}{\epsilon^{2} \alpha_{0}^{2} l\left(\alpha_{0}-(d-1) \epsilon / d\right)} \\
& \geq 1-\frac{d^{3} \alpha_{i}}{\epsilon^{3} \alpha_{0}^{2} l},
\end{aligned}
$$

since $\alpha_{0}>\epsilon$.
This gives us

$$
\begin{aligned}
\operatorname{Pr}\left[T_{0} \wedge \cdots \wedge T_{d-1}\right] & =\prod_{i=1}^{d-1} \operatorname{Pr}\left[T_{i} \mid T_{i+1} \wedge \cdots \wedge T_{d-1}\right] \\
& \geq \prod_{i=1}^{d-1}\left(1-\frac{d^{3} \alpha_{i}}{\epsilon^{3} \alpha_{0}^{2} l}\right) \\
& \geq 1-\frac{d^{3}\left(1-\alpha_{0}\right)}{\alpha_{0}^{2} \epsilon^{3} l}
\end{aligned}
$$

So with $\beta=1-\frac{d^{3}\left(1-\alpha_{0}\right)}{\alpha_{0}^{2} \epsilon^{3} l}$, the result follows.

## Chapter 3

## Provable security

We shall give a brief introduction to the setting in which we do our work. The goal of provable security is to provide a reduction from the security of the cryptosystem to some well-studied problem (the underlying problem). By itself, this says nothing about the security of the cryptosystem. But if we make the assumption that there are no algorithms that solve the underlying problem with a significant probability and at a cost less than some bound, the reduction then says something about the minimal cost for breaking the cryptosystem.

The main goal of this chapter is to explain what "breaking the cryptosystem" may mean, that is, what the security requirements of public key cryptosystems are, and what kind of attacks we allow. We shall also discuss the different models in which we obtain our security proofs, and the assumptions underlying those models.

### 3.1 Public key cryptosystems

The following definition of a public key cryptosystem is suitable for our purposes.
Definition. A public key cryptosystem $\Pi$ consists of three algorithms, $\mathcal{K}, \mathcal{E}$ and $\mathcal{D}$. The key generation algorithm $\mathcal{K}$ is an efficient algorithm that takes a security parameter as input and outputs a public key $p k$ and private key sk. The public key specifies among other things a finite set of possible messages and a finite set of possible ciphertexts, denoted by $p k_{M}$ and $p k_{C}$.

The encryption algorithm $\mathcal{E}$ is an efficient algorithm that takes a public key and a message as input and outputs a ciphertext. The decryption algorithm $\mathcal{D}$ is a deterministic polynomial-time algorithm that takes a private key and a ciphertext as input and outputs either a message or the special symbol $\perp$. We require that for any public-private key pair ( $p k, s k$ ) and any message $m$ in the set of messages specified by $p k, \mathcal{D}(s k, \mathcal{E}(p k, m))=m$.

Note that the security parameter input to $\mathcal{K}$ should be input as a string of 1 's of length $\tau$ (denoted by $1^{\tau}$ ), not as (say) a binary encoding. Since we require the
algorithm to be efficient, the number of steps must be bounded by a polynomial in the input length. In the former case, this would be a polynomial in $\tau$. In the latter case, it would be a polynomial in $\log _{2} \tau$, which is wrong. We will henceforth ignore this slightly subtle point, and use the notation $\mathcal{K}(\tau)$ as a short-hand for $\mathcal{K}\left(1^{\tau}\right)$.

We shall also be interested in public key cryptosystems with a certain property.

Definition. Let $\Pi=(\mathcal{K}, \mathcal{E}, \mathcal{D})$ be a public key cryptosystem and let $(p k, s k)$ be a public-private key pair for $\Pi$. Let $*: p k_{M} \times p k_{M} \rightarrow p k_{M}$ and $*^{\prime}: p k_{C} \times p k_{C} \rightarrow p k_{C}$ be binary operations such that for all $c, c^{\prime} \in p k_{C}$, we have that

$$
\mathcal{D}\left(c *^{\prime} c^{\prime}\right)=\mathcal{D}(c) * \mathcal{D}\left(c^{\prime}\right)
$$

We then say that $\Pi$ is homomorphic with respect to $*$ and $*^{\prime}$.

### 3.2 Adversary models

In a public key cryptosystem, the adversary obviously knows the public key. This means that he can always mount a chosen plaintext attack, wherein he can encrypt any message of his choice with the given public key.

We can assume that the adversary observes exactly one ciphertext (the challenge ciphertext) that he does not know the decryption of. Observing more than one unknown ciphertext does not help him, since he can create ciphertexts himself, and then forget about their decryptions.

We model the adversary as two separate stages. First, the adversary receives the public key and may do some arbitrary computations. He then outputs some state information that is passed on to the second stage, and possibly some information that influences the challenge ciphertext. This is often called the "find"-stage.

After the first stage is complete, the challenge ciphertext is created outside of the adversary's control. It may, depending on the attack, be influenced by the output of the first stage, but obviously not determined by it.

The second stage receives the public key, the state information output by the first stage and the challenge ciphertext. It tries to determine some information about the challenge ciphertext. This is often called the "guess"-stage.

The adversary $A$ is then a pair of algorithms $\left(A_{1}, A_{2}\right)$, where $A_{1}$ carries out the first stage, and $A_{2}$ carries out the second.

It is conceivable that the adversary can do more than just computations based on the public key and the ciphertext. If the adversary can somehow obtain decryptions of ciphertexts he has created, he may be able to learn something about the secret key. If the attacker can modify a ciphertext and somehow get some information about the decryption of the modified ciphertext, it is possible that he can recover some information about the decryption of the original ciphertext, or the secret key.

This suggests that the adversary should be allowed to do more than just computations. To model this, we allow the adversary to decrypt ciphertexts using the secret key. One variant is to give the adversary access to the decryption oracle during the first stage [30]. This is called a non-adaptive chosen ciphertext attack.

While this may in certain situations be a realistic scenario, it is usually unrealistic to assume that the adversary does not have any additional powers during the second stage. The adversary's task is obviously trivial if we give him access to a decryption oracle in the second stage. The solution [36] is to give him access to a restricted decryption oracle, that refuses to decrypt the challenge ciphertext, but decrypts any other ciphertext. This is called an adaptive chosen ciphertext attack, or simply a chosen ciphertext attack.

In practice, this seems to give the adversary unreasonable advantages. But it makes sense, since if a cryptosystem is secure against a (non-adaptive) chosen ciphertext adversary, then it will certainly be secure against less powerful, realistic adversaries.

### 3.3 Security goals

The decryption problem for $\Pi$ and a fixed public-private key pair $(p k, s k)$ is the set of possible ciphertexts output by $\mathcal{E}$ when the message is sampled uniformly at random from $p k_{M}$. The adversary is given the public key $p k$ and the ciphertext $c$. The answer to each instance is the decryption of the ciphertext.

Definition. A public key cryptosystem $\Pi$ is one-way if the corresponding decryption problem is hard.

In many practical situations, one-way is not sufficient to guarantee security. Goldwasser and Micali [22] first proposed the notion of semantic security. This is a variant of Shannon's perfect secrecy $[38]$ for adversaries with bounded computational power. The idea is that it should be difficult to decrypt a ciphertext even when the message space is restricted, and that it should be difficult to recover any information about the decryption of the ciphertext, beyond what is a-priori known from the message distribution.

Let $\Pi=(\mathcal{K}, \mathcal{E}, \mathcal{D})$ be a public key cryptosystem. We shall allow the adversary to influence the choice of message. We do this by allowing the adversary's "find"stage to output an arbitrary function $f: p k_{M} \rightarrow\{0,1\}$ along with a probability space $X$ on $p k_{M}$ such that

$$
|\operatorname{Pr}[f(x)=0 \mid x \leftarrow X]-1 / 2|
$$

is negligible. The message $m$ to be encrypted is the sampled from $X$, and the adversary's "guess"-stage is given the challenge ciphertext. Its task is to guess $f(m)$.

We assume that there are efficient algorithms available for sampling $X$ and computing $f$.

We use the following experiment to measure the adversary's advantage.

## Experiment 1.

Input: $\Pi, A=\left(A_{1}, A_{2}\right)$, $\tau$.

1. $(p k, s k) \leftarrow \mathcal{K}(\tau)$.
2. $(X, f, o) \leftarrow A_{1}(p k)$.
3. $m \leftarrow X, c \leftarrow \mathcal{E}(p k, m)$.
4. $b \leftarrow A_{2}(p k, c, o)$.
5. If $f(m)=b$, output 1 , otherwise output 0 .

Output: 0 or 1 .
Here, $o$ is some state information that $A_{1}$ wants to pass on to $A_{2}$. The adversary's advantage is then defined to be

$$
\operatorname{Adv}_{A}(\tau)=|\operatorname{Pr}[\operatorname{Exp} \mathbf{1}(\Pi, A, \tau)=1]-1 / 2|
$$

Definition. We say that $\Pi$ is semantically secure if any efficient adversary as above has negligible advantage.

We note that some definitions of semantic security require $X$ to be uniformly distributed. Also note that semantic security is not a problem in the sense of Section 2.2, since the problem solver participates in an arbitrary way in sampling the challenge ciphertext and the answer.

It is quite obvious that any system that is semantically secure is also one-way.
Working with semantic security is awkward, but there are simpler notions available. Suppose the adversary's "find"-stage outputs a probability space $X$ given by

$$
\operatorname{Pr}[X=m]= \begin{cases}\frac{1}{2} & m=m_{0} \text { or } m=m_{1} \\ 0 & \text { otherwise }\end{cases}
$$

where $m_{0}, m_{1}$ are two distinct messages in $p k_{M}$, and the function $f$ is defined to be $f\left(m_{b}\right)=b$, and otherwise arbitrary.

The adversary's "guess"-stage must then decide which of the two messages the challenge ciphertext decrypts to.

Definition. We say that $\Pi$ has indistinguishable encryptions if any efficient adversary as described above has probability negligibly different from $1 / 2$ in guessing which message was encrypted.

Indistinguishable encryptions is often called polynomial security in the literature.

It is quite obvious that if $\Pi$ is semantically secure, then $\Pi$ has indistinguishable encryptions. The converse is also true [22].

Theorem 3.1. Let $\Pi$ be a public key cryptosystem. If $\Pi$ has indistinguishable encryptions, then $\Pi$ is semantically secure.

Proof. Suppose we have an adversary $A=\left(A_{1}, A_{2}\right)$ against semantic security with significant advantage. We shall turn it into an adversary $A^{\prime}=\left(A_{1}^{\prime}, A_{2}^{\prime}\right)$ against indistinguishable encryptions. $A_{1}^{\prime}$ runs $A_{1}$ and receives $(X, f)$. It samples $m_{0}, m_{1}$ from $X$ such that $f\left(m_{b}\right)=b$. This can be done efficiently with arbitrarily small probability of failure.

When $A_{2}^{\prime}$ receives the challenge ciphertext $c$, which is an encryption of $m_{b}$, it runs $A_{2}$ and receives its guess $b^{\prime}$ for $f\left(m_{b}\right)=b$. Since the message $m_{b}$ has been sampled according to $X$, the input to $A_{2}$ is exactly as it expects. It therefore guesses $f\left(m_{b}\right)=b$ correctly with a significant advantage, hence $A_{2}^{\prime}$ guesses $b$ correctly with significant advantage.

A second alternative security goal is the following. The "find"-stage outputs the probability space $X$ given by

$$
\operatorname{Pr}[X=m]= \begin{cases}\frac{1}{2} & m=m_{0} \\ \frac{1}{\left|p k_{M}\right|-1} & \text { otherwise }\end{cases}
$$

where $m_{0} \in p k_{M}$, and the function $f$ is defined to be $f\left(m_{0}\right)=1$ and $f(m)=0$, when $m \neq m_{0}$.

The adversary's "guess"-stage must then decide if the challenge ciphertext decrypts to $m_{0}$ or not.

Definition. We say that $\Pi$ has encryptions indistinguishable from random if any efficient adversary as described above has probability negligibly different from $1 / 2$ in guessing if the message $m_{0}$ was encrypted.

It is quite obvious that if $\Pi$ is semantically secure, then $\Pi$ has encryptions indistinguishable from random. The converse is also true.

Theorem 3.2. Let $\Pi$ be a public key cryptosystem. If $\Pi$ has encryptions indistinguishable from random, then $\Pi$ is semantically secure.

Proof. Suppose we have an adversary $A=\left(A_{1}, A_{2}\right)$ against indistinguishable encryptions with significant advantage $\epsilon$. We shall turn it into an adversary $A^{\prime}=\left(A_{1}^{\prime}, A_{2}^{\prime}\right)$ against encryptions indistinguishable from random. The result will then follow by Theorem 3.1 if $A^{\prime}$ has significant advantage.
$A_{1}^{\prime}$ runs $A_{1}$ and receives $m_{0}, m_{1} \in p k_{M}$. It flips a coin $b$ and outputs $m_{b}$. $A_{2}^{\prime}$ runs $A_{2}$ and receives $b^{\prime}$. If $b^{\prime}=b, A_{2}^{\prime}$ outputs 1 , otherwise it outputs 0 .

The ciphertext given to $A_{2}^{\prime}$ is either an encryption of $m_{b}$, or of some other random message. If the challenge ciphertext decrypts to $m_{b}$, it is clear that $A_{2}$ guesses $b$ correctly with probability $1 / 2+\epsilon$.

If the challenge ciphertext does not decrypt to $m_{b}$, we have two possibilities. If it decrypts to $m_{1-b}$ (which happens with probability $\left.1 /\left(\left|p k_{M}\right|-1\right)\right), A_{2}$ guesses
$b$ correctly with probability $1 / 2-\epsilon$. If the challenge ciphertext does not decrypt to $m_{1-b}$ (which happens with probability $\left(\left|p k_{M}\right|-2\right) /\left(\left|p k_{M}\right|-1\right), A_{2}$ has no information about $b$ and guesses correctly with probability $1 / 2$.

Summing up, we get

$$
\begin{aligned}
\operatorname{Adv}_{A^{\prime}} & =\frac{1}{2}\left|\operatorname{Pr}\left[A_{2}^{\prime}=1 \mid \mathcal{D}(c)=m_{b}\right]-\operatorname{Pr}\left[A_{2}^{\prime}=1 \mid \mathcal{D}(c) \neq m_{b}\right]\right| \\
& =\frac{1}{2}\left|\frac{1}{2}+\epsilon-\left(\frac{1}{2}-\epsilon\right) \frac{1}{\left|p k_{M}\right|-1}-\frac{1}{2} \frac{\left|p k_{M}\right|-2}{\left|p k_{M}\right|-1}\right| \\
& =\frac{1}{2} \epsilon \frac{\left|p k_{M}\right|}{\left|p k_{M}\right|-1},
\end{aligned}
$$

which concludes the proof.
From the proof, we see that if $\left|p k_{M}\right|=2$, the two notions indistinguishable encryptions and encryptions indistinguishable from random coincide, as they should.

A different security goal is non-malleability. Briefly, it says that it is difficult to modify a ciphertext in such a way that the decryption of the modified ciphertext is predictably related to the original ciphertext. We shall not work with nonmalleability, and we only remark that no homomorphic cryptosystem can be non-malleable.

### 3.4 Security notions

The security goals are independent of the adversary models, so we can mix them freely to form security notions [3]. It is now interesting to ask if a cryptosystem that satisfies some security notion also satisfies other security notions.

It is clear that if a cryptosystem achieves some security goal against some adversary class, then it achieves the same security goal against weaker adversary classes. It is also easy to show that any cryptosystem that achieves nonmalleability against some adversary class also achieves semantic security against the same adversary class.

For a general-purpose public key cryptosystem, one wants both semantic security and non-malleability against adaptive chosen ciphertext adversaries. This ensures that an adversary cannot learn anything about the message encrypted, and he cannot tamper with the ciphertext (the ciphertext "envelope" containing the message is both opaque and tamper-proof).

It has been shown [3] that any cryptosystem that is semantically secure against chosen ciphertext adversaries also achieves non-malleability against chosen ciphertext adversaries. Since non-malleability is somewhat difficult to work with, this result simplifies our work.

Since no homomorphic public key cryptosystem can be non-malleable, semantic security against non-adaptive chosen ciphertext adversaries is the strongest security notion a homomorphic cryptosystem can achieve.

We shall use semantic security as a shorthand for semantic security against chosen plaintext attack, and abbreviate this by IND. We shall say that a cryptosystem is secure against (non-adaptive) chosen ciphertext attacks if it is semantically secure against (non-adaptive) chosen ciphertext attacks, and abbreviate this by CCA1 and CCA2 ${ }^{1}$.

### 3.5 Key encapsulation methods

This presentation is based on [11].
While sending short messages has many applications, quite often one wants to send a long message using public-key methods. Since public key cryptosystems are quite slow, sending long messages would take a long time. Usually, symmetric (or private key) cryptosystems are much faster than public key cryptosystems. Hence, it could be interesting to combine a symmetric cryptosystem with a public key cryptosystem.

One could do this by using the public key cryptosystem to encrypt a key for a symmetric cryptosystem, and then encrypt the message using the symmetric cryptosystem. But the keys used for a symmetric cryptosystem are in general chosen uniformly at random from some key space. Because of this, there may be better approaches to designing public key cryptosystems for long messages.

A (one-time) symmetric cryptosystem $\Sigma$ consists of two algorithms $\mathcal{S E}$ and $\mathcal{S D}$, a set of symmetric keys $\Sigma_{K}$, a set of messages $\Sigma_{M}$ and a set of ciphertexts $\Sigma_{C}$. The symmetric encryption algorithm $\mathcal{S E}$ takes as input a symmetric key and a message, and outputs a ciphertext. The symmetric decryption algorithm $\mathcal{S D}$ is deterministic, takes as input a symmetric key and a ciphertext, and outputs a message or the special symbol $\perp$. We require that for all message $m \in \Sigma_{M}$ and $k \in \Sigma_{K}, \operatorname{Pr}[\mathcal{S D}(k, \mathcal{S E}(k, m))=m]=1$.

Note that the length of distinct messages may be different, and the length of the ciphertexts may depend on the length of the message.

We say that a symmetric cryptosystem is secure against passive attacks if any adversary that first selects two messages $m_{0}$ and $m_{1}$ (of equal length), receives an encryption of either message under some secret key, and then tries to guess which message was encrypted, guesses correctly with negligible advantage. Note that the adversary only receives a single ciphertext under any one key, hence "one-time".

A symmetric cryptosystem is secure against active attacks if the advantage of any adversary that is allowed to decrypt any ciphertext under the secret key (except, of course, the challenge ciphertext), is negligible.

Typical symmetric cryptosystems include the one-time pad, block ciphers in a secure mode of encryption, and general stream ciphers. To protect against active

[^1]attacks, a message authentication code (MAC) is usually added, following the encrypt-then-authenticate paradigm.

Definition. A key encapsulation method (KEM) $\Lambda$ consists of three algorithms, $\mathcal{K}, \mathcal{E}$ and $\mathcal{D} . \mathcal{K}$ is an efficient algorithm that takes a security parameter as input and outputs a public key $p k$ and private key sk. The public key specifies among other things a finite set of possible symmetric keys and a finite set of possible ciphertexts, denoted by $p k_{K}$ and $p k_{C}$.
$\mathcal{E}$ is an efficient algorithm that takes a public key as input and outputs a symmetric key and a ciphertext encapsulating the symmetric key. $\mathcal{D}$ is a deterministic polynomial-time algorithm that takes a private key and a ciphertext as input and outputs either a symmetric key or the special symbol $\perp$. We require that for any public-private key pair $(p k, s k), \operatorname{Pr}[\mathcal{D}(s k, c)=k \mid(k, c) \leftarrow \mathcal{E}(p k)]=1$.

The intuitive notion is that the key encapsulation method is secure if it is impossible to distinguish the encapsulated symmetric key from a symmetric key chosen uniformly at random. We use the following experiment to measure the advantage of an adversary.

## Experiment 2.

Input: $\Lambda=(\mathcal{K}, \mathcal{E}, \mathcal{D}), A, \tau$.

1. $(p k, s k) \leftarrow \mathcal{K}(\tau)$.
2. $(x, c) \leftarrow \mathcal{E}(p k)$.
3. $b \leftarrow\{0,1\}$.
4. If $b=1$, then $z \leftarrow x$, otherwise $z \leftarrow p k_{K}$.
5. $b^{\prime} \leftarrow A(p k, c, z)$.
6. If $b=b^{\prime}$, output 1 , otherwise output 0 .

Output: 0 or 1 .
We define the advantage of an adversary $A$ to be

$$
\operatorname{Adv}_{A}(\tau)=|\operatorname{Pr}[\operatorname{Exp} 2(\Lambda, A, \tau)=1]-1 / 2| .
$$

In a chosen ciphertext attack, we give the adversary access to a restricted decryption oracle. (We do not distinguish between adaptive and non-adaptive chosen ciphertext attacks.)

We say that a key encapsulation method is semantically secure if any efficient passive adversary has negligible advantage, and it is secure against chosen ciphertext attacks if any efficient chosen ciphertext adversary has negligible advantage.

A hybrid cryptosystem is a pair $(\Lambda, \Sigma)$. The key generation algorithm takes as input a security parameter, and outputs a public key for $\Lambda$ such that $p k_{K}=\Sigma_{K}$. The encryption and decryption algorithms are given in Figure 3.1. Depending on the symmetric cryptosystem, the range of the security parameter for the key generation algorithm may be restricted.

The security goals and adversary models for hybrid cryptosystems are as for public key cryptosystems. It is possible to prove the following theorem.

## Encryption.

Input: $p k, m \in \Sigma_{M}$.

1. $(k, c) \leftarrow \mathcal{E}(p k)$.
2. $c^{\prime} \leftarrow \mathcal{S E}(k, m)$.
3. Output $\left(c, c^{\prime}\right)$.

Output: A hybrid ciphertext in $p k_{C} \times \Sigma_{C}$.

## Decryption.

Input: $s k,\left(c, c^{\prime}\right) \in p k_{C} \times \Sigma_{C}$.

1. $k \leftarrow \mathcal{D}(s k, c)$.
2. If $k=\perp$, output $\perp$.
3. $m \leftarrow \mathcal{S D}\left(k, c^{\prime}\right)$.
4. Output $m$.

Output: A message $m \in \Sigma_{M}$ or the special symbol $\perp$.

Figure 3.1: Encryption and decryption algorithms for hybrid cryptosystems.

Theorem 3.3. The hybrid cryptosystem $(\Lambda, \Sigma)$ is semantically secure if $\Lambda$ is semantically secure and $\Sigma$ is secure against passive attacks. ( $\Lambda, \Sigma$ ) is secure against chosen ciphertext attacks if $\Lambda$ is secure against chosen ciphertext attacks and $\Sigma$ is secure against active attacks.

### 3.6 Models for provable security

In general, we want to prove the security of a cryptosystem using as few assumptions as possible. If we place no limitations on the adversary, except that imposed by the computational model and certain bounds on computational power, we are working in the real world, and we call this the standard model.

### 3.6.1 Random oracle model

A hash function is a function that takes as input arbitrary binary strings and outputs a fixed-length bit string. We are interested in a special class of hash functions with desirable cryptographic properties.

Definition. Let $h:\{0,1\}^{*} \rightarrow\{0,1\}^{s}$ be a hash function. If $x_{0}$ and $x_{1}$ are distinct bit strings such that $h\left(x_{0}\right)=h\left(x_{1}\right)$, we say that the pair $\left(x_{0}, x_{1}\right)$ is a collision for $h$. A collision resistant hash function is a hash function such that the problem of finding collisions is hard.

There are families of hash functions that are provably collision resistant, but they tend to be very inefficient and are not interesting in practice.

Several very efficient hash functions have been proposed, that seem to be collision resistant, such as the SHA-i [19]. The interesting thing about these hash functions is that they seem to behave as random functions. Informally, we can say that the only way to get any information about the value of a random function at some point is to evaluate the function.

It seems difficult to capture this notion in the standard model, so the random oracle model was formulated [5]. The idea is to prove security using random functions, and then replace the random functions with hash functions like SHA-i.

Let $S$ and $S^{\prime}$ be two sets, and let $\mathcal{H}$ be the set of all functions from $S$ to $S^{\prime}$. A random oracle from $S$ to $S^{\prime}$ is an oracle that is made available to all algorithms and oracles that participate in an experiment. At the start of the experiment, the oracle is initialised by choosing a function $h$ from $\mathcal{H}$ uniformly at random. Whenever the random oracle receives a query $x$, it computes $h(x)$ and returns the result. The other algorithms and oracles participating in the experiment can only get information about $h$ by querying the oracle.

An alternative formulation that is often convenient, is a description of a random oracle in terms of an algorithm. The oracle algorithm accepts as queries elements of the set $S$. For the first query $x$ the oracle receives, the oracle samples an element $y$ of $S^{\prime}$ uniformly at random. It then stores the pair $(x, y)$ in a sorted list, and returns the answer $y$.

For each subsequent query $x^{\prime}$, the oracle checks if it has a pair $\left(x^{\prime}, y^{\prime}\right)$ in its list. If it does have such a pair, it returns the answer $y^{\prime}$. If it does not have such a pair, it samples an element $y^{\prime}$ of $S^{\prime}$ uniformly at random, stores the pair $\left(x^{\prime}, y^{\prime}\right)$ in the sorted list, and returns the answer $y^{\prime}$.

It is quite clear that these two formulations are equivalent.
While the random oracle model is quite popular, a proof of security in the random oracle model does not imply that the cryptosystem is secure when the random oracle is replaced with a real hash function. While the security certainly depends on what hash function is used, there are also schemes [2] that are provably secure in the random oracle model, but are insecure no matter what hash function is used.

Therefore, a proof of security in the random oracle model should only be taken as heuristic evidence for the security of the system, and great care should be taken when instantiating the system.

### 3.6.2 Generic model

It is interesting to study algorithms that operate on arbitrary finite abelian groups. These are often called generic algorithms. Abstractly, a finite abelian group is isomorphic to a product of different $\mathbb{Z}_{p^{r}}$, but for some groups, this isomorphism seems to be difficult to compute. Shoup [39] showed that no generic algorithm can compute that isomorphism efficiently.

A consequence is that when the set $\Gamma_{\tau}$ is the set of abstract groups of prime order $p$, where $p$ is a $\tau$-bit prime and $\tau>1$ and $\Gamma=\left\{\Gamma_{\tau}\right\}$, then the discrete logarithm problem $\mathcal{D} \mathcal{L}_{\Gamma}$ is provably hard. This shows that there can be no generic attacks on the discrete logarithm problem, but it does not say anything about how difficult the discrete logarithm problem is for real families of groups.

The generic model works as follows. Let $\mathbb{Z}_{n}$ be the group structure, and let $\mathcal{H}$ be the set of 1-1 functions from $\mathbb{Z}_{n}$ into a finite set $S$, where $|S| \geq n$. At the start of the experiment, a group oracle is initialised with a function $f \in \mathcal{H}$ chosen
uniformly at random.
The group oracle accepts three types of queries. The first query type consists of the special symbol $\perp$, and the oracle responds with $f(0)$ and $f(1)$. The second query type consists of a single element $s \in S$. If $s \notin f\left(\mathbb{Z}_{n}\right)$, then the group oracle answers $\perp$, otherwise it returns $f\left(-f^{-1}(s)\right)$ (note that $f^{-1}(s)$ is well-defined for $\left.s \in f\left(\mathbb{Z}_{n}\right)\right)$. The third query type consists of a pair $\left(s_{1}, s_{2}\right) \in S \times S$. If $s_{1} \notin f\left(\mathbb{Z}_{n}\right)$ or $s_{2} \notin f\left(\mathbb{Z}_{n}\right)$, then the answer is $\perp$, otherwise the answer is $f\left(f^{-1}\left(s_{1}\right)+f^{-1}\left(s_{2}\right)\right)$.

Every algorithm and oracle in the experiment is given access to this group oracle.

As for the random oracle model, we can also describe the group oracle in terms of an algorithm.

The algorithm keeps two lists: one list of pairs $\{(s, l)\} \subseteq S \times \mathbb{Z}_{n}$, sorted on both coordinates, and one list of "forbidden" elements $\{s\} \subseteq S$. The first list will be a partial definition of a 1-1 function $f: \mathbb{Z}_{n} \rightarrow S$, so we call it the $f$-list. The second list will contain a partial list of $S \backslash f\left(\mathbb{Z}_{n}\right)$.

If $s \in S$ is received in a query, and $s$ is in the forbidden list, $\perp$ is returned immediately.

Whenever $s \in S$ is received in a query, there is no pair $(s, l)$ in the $f$-list, and $s$ is not in the "forbidden" list, the algorithm flips a biased coin. The coin is 0 with probability $(n-N) /(|S|-N)$, where $N$ is the number of elements in the $f$-list. If the coin flip gives $0, s$ is stored in the "forbidden" list and $\perp$ is returned. If the coin flip gives $1, l$ is sampled uniformly at random from $\mathbb{Z}_{n}$ such that there is no pair $\left(s^{\prime}, l\right)$ in the $f$-list, and then $(s, l)$ is added to the $f$-list, and the query proceeds as usual.

Otherwise, the query is processed as usual, and the algorithm should return $f(l)$ for some $l \in \mathbb{Z}_{n}$. If there is no pair $\left(s^{\prime}, l\right)$ in the $f$-list, then $s^{\prime}$ is sampled uniformly from $S$ such that there is no pair $\left(s^{\prime}, l^{\prime}\right)$ in the $f$-list, and such that $s^{\prime}$ is not in the forbidden list. Then $\left(s^{\prime}, l\right)$ is added to the $f$-list, and $s^{\prime}$ is returned.

It is clear that the two descriptions of the group oracle are equivalent.
Quite often, cryptographic systems are formulated in terms of generic abelian groups. In this case, it may make sense to prove that the system is secure against generic algorithms. As for the random oracle model, this does not actually imply that a concrete instantiation of the generic system is secure. But it does provide some heuristic evidence (albeit weaker than a proof in the random oracle model) for the security of concrete instantiations.

### 3.7 Plaintext aware encryption

We shall discuss one technique that is often used to prove security. The basic idea is that if the adversary is unable to create ciphertexts without knowing their decryption, a decryption oracle is useless and (non-adaptive) chosen ciphertext attacks reduce to chosen plaintext attacks.

Let $\Pi=(\mathcal{K}, \mathcal{E}, \mathcal{D})$ be a public key cryptosystem. A forger for $\Pi$ is an algorithm $B$ that takes as input a public key $p k$ and outputs a ciphertext $c$. We give $B$ access
to a special encryption oracle $\mathcal{E}_{p k}$. When it is first queried with a probability space $X$ on $p k_{M}$, it samples a message $m_{0}$ according to $X$, a ciphertext $c_{0}$ according to $\mathcal{E}(p k, m)$, and returns the answer $c_{0}$. Any subsequent query is answered with the special symbol $\perp$.

It is quite clear that we can modify any such forger to also output $c_{0}$ in addition to $c$.

A knowledge extractor for $\Pi$ is an algorithm $D$ that on input of a public key $p k$, a ciphertext $c_{0}$, a trace and a challenge ciphertext $c$, outputs a message $m \in p k_{M}$ or the special symbol $\perp$. The idea is that the trace should be enough for a knowledge extractor to decrypt the ciphertext.

Removing the oracle $\mathcal{E}_{p k}$ and the ciphertext $c_{0}$ from the above gives us definitions for weak forger and weak knowledge extractor.

We use the following experiment to measure the success probability of a knowledge extractor. (The weak variant is obvious.)

## Experiment 3.

Input: $\Pi=(\mathcal{K}, \mathcal{E}, \mathcal{D}), B^{\prime}, D, \tau$.

1. $(p k, s k) \leftarrow \mathcal{K}(\tau)$.
2. Initialise the $\mathcal{E}_{p k}$ oracle.
3. $\left(c, c_{0}, t r\right) \leftarrow B^{\prime \mathcal{E}_{p k}}(p k)$.
4. If $c=c_{0}$, output 0 and stop.
5. $m \leftarrow D\left(p k, c_{0}, t r, c\right)$.
6. $m^{\prime} \leftarrow \mathcal{D}(s k, c)$.
7. If $m=m^{\prime}$, output 1 , otherwise output 0 .

Output: 0 or 1 .
Now we can say exactly what it means that a public key cryptosystem is plaintext aware.

Definition. A public key cryptosystem $\Pi$ is (weakly) plaintext aware if there exists a (weak) knowledge extractor $D$ such that for any (weak) forger $B$ for $\Pi$, there exists a trace-variant $B^{\prime}$ of $B$ such that

$$
\operatorname{Pr}\left[\operatorname{Exp} 3\left(\Pi, B^{\prime}, D, \tau\right)=0\right]
$$

is negligible in $\tau$.
The following theorem can then be proven (as in [3]).
Theorem 3.4. Let $\Pi$ be a public key cryptosystem that is semantically secure and (weakly) plaintext aware. Then $\Pi$ is secure against (non-adaptive) chosen ciphertext attacks.

The main idea of the proof is that every time the adversary queries its decryption oracle, it is in effect a (weak) forger. This means that there exists a trace-variant that makes the same query, but where the query includes the trace. The decryption oracle can then use the trace along with the knowledge extractor to do the decryption.

The requirement that $\Pi$ is semantically secure is often included in the definition of plaintext awareness.

Plaintext awareness is especially interesting in the random oracle model. It is quite obvious that we can modify any algorithm in the random oracle model to output, in addition to its normal output, a complete list of any queries it has made to any random oracles it has access to. So in the random oracle model, the list of random oracle queries is a very interesting trace for a forger to output.

## Chapter 4

## Subgroup membership problems

As we have seen, the strongest security goals are formulated in terms of indistinguishability. There is a class of general mathematical problems, called subgroup membership problems, that has a very natural indistinguishability-property.

### 4.1 Subgroup membership problems

Let $\Gamma_{\tau}$ be a probability space on a set of pairs $(G, K)$, where $G$ is a finite abelian group and $K$ is a subgroup of $G$. For each pair $(G, K)$, define a probability space $X_{G}$ on $G$ such that $X_{G}$ when restricted to $K$ or $G \backslash K$ is uniformly distributed, and

$$
\operatorname{Pr}\left[x \in K \mid x \leftarrow X_{G}\right]=\operatorname{Pr}\left[x \in G \backslash K \mid x \leftarrow X_{G}\right]=\frac{1}{2}
$$

Let $\Gamma=\left\{\Gamma_{\tau}\right\}$.
We assume that an efficient algorithm is available for $\Gamma$ that samples $K$ and $G \backslash K$ (almost) uniformly at random, and we denote sampling $x$ from the resulting probability space by $x \approx K$ and $x \approx G \backslash K$.
Definition. The subgroup membership problem $\mathcal{S M}_{\Gamma}$ derived from the family $\Gamma$ described above is the following problem: The instances are triples $(G, K, x)$, where $(G, K)$ is sampled from $\Gamma_{\tau}$ and $x$ is sampled according to $X_{G}$. The answer to the instance $(G, K, x)$ is 1 if $x \in K$, otherwise 0 .

We use the word distinguisher for an algorithm that solves a subgroup membership problem.

We denote by $\mathcal{S M}_{(G, K)}$ the subproblem of $\mathcal{S} \mathcal{M}_{\Gamma}$ where the instances are restricted to triples $(G, K, x)$, with $x \in G$. It is clear that if we can show that something holds for an arbitrary subproblem $\mathcal{S M}_{(G, K)}$, then it also holds for $\mathcal{S M}_{\Gamma}$. (Note that the subgroup membership problem is equivalent to deciding if the residue of $x$ is the neutral element in the factor group $G / K$.)

It is easy to see that for $\mathcal{S M}_{(G, K)}$, we have that

$$
\begin{aligned}
& \left.\operatorname{Adv}_{A}^{\mathcal{S} \mathcal{M}_{(G, K)}}=\frac{1}{2} \right\rvert\, \operatorname{Pr}[A(G, K, x)=1 \mid x \leftarrow K]- \\
& \operatorname{Pr}[A(G, K, x)=1 \mid x \leftarrow G \backslash K] \mid .
\end{aligned}
$$

Proposition 4.1. Suppose $\mathcal{S M}_{(G, K)}$ is a subgroup membership problem, that the factor group $G / K$ is cyclic and $|G / K|$ has no small prime factors, or that $|G / K|$ has known, prime order. Then $\mathcal{S M}_{(G, K)}$ is self-reducible.
Proof. If $|G / K|$ has no small prime factors, let $t_{0}$ be such that $|G / K|<2^{t_{0}}$. Let $U$ be uniformly distributed on the set $\left\{0, \ldots, 2^{t_{0}+t}-1\right\}$, for some $t>0$. By Lemma 2.3, $U \bmod |G|$ is $2^{-t}$-close to uniformly distributed on $\{0, \ldots,|G|-1\}$.

Alternatively, if $|G / K|$ has known, prime order $\ell$, then let $U$ be uniformly distributed on $\{1, \ldots, \ell-1\}$.

Let $f_{u}$ be the map $x \mapsto x^{u}$. Sampling $u$ from $U$ samples $f_{u}$ (almost) uniformly from the set of automorphism on $G / K$. (With small probability, it could be something other than an automorphism.)

The algorithm takes as input $x \in G$. It samples $u$ from $U$ and $x^{\prime}$ from $K$, and computes $x^{\prime \prime}=x^{\prime} x^{u}$, then it outputs $x^{\prime \prime}$ and the identity function.

If the input $x \in K$, then it is clear that $x^{\prime \prime} \in K$ and its distribution is (almost) uniform and independent of $x$.

If $x \in G \backslash K$, then except with small probability, the residue class of $x$ in $G / K$ is a generator, and so $x^{U}$ is (almost) uniformly distributed in $G / K$. It then follows that $x^{U} x^{\prime}$ is (almost) uniformly distributed in $G \backslash K$ and it is independent of $x$.

Proposition 4.2. Let $G$ be a group and let $K^{\prime} \subseteq K \subseteq G$ be proper subgroups. If $A$ can distinguish $K^{\prime}$ from $G$ with advantage $\epsilon$, then either $A$ can distinguish $K^{\prime}$ from $K$ with advantage greater than or equal to $\epsilon / 2$, or it can distinguish $K$ from $G$ with advantage greater than or equal to $\epsilon / 3$. Alternatively,

$$
\operatorname{Adv}_{A}^{\mathcal{S} \mathcal{M}_{\left(G, K^{\prime}\right)}} \leq \operatorname{Adv}_{A}^{\mathcal{S} \mathcal{M}_{(G, K)}}+\frac{3}{2} \operatorname{Adv}_{A}^{\mathcal{S} \mathcal{M}_{\left(K, K^{\prime}\right)}}
$$

Proof. Set $X_{0}=K^{\prime}, X_{1}=K \backslash K^{\prime}$ and $X_{2}=G \backslash K$. Let $w t(X)=\operatorname{Pr}\left[A\left(G, K^{\prime}, x\right)=\right.$ $1 \mid x \in X]$. We may without loss of generality assume that $\epsilon>0$ and that

$$
\begin{aligned}
2 \epsilon & =w t\left(X_{1} \cup X_{2}\right)-w t\left(X_{0}\right) \\
& =\frac{\left|X_{1}\right|}{\left|X_{1}\right|+\left|X_{2}\right|} w t\left(X_{1}\right)+\frac{\left|X_{2}\right|}{\left|X_{1}\right|-\left|X_{2}\right|} w t\left(X_{2}\right)-w t\left(X_{0}\right) .
\end{aligned}
$$

Setting $2 \delta_{0}=w t\left(X_{1}\right)-w t\left(X_{0}\right)$ and $2 \delta_{1}=w t\left(X_{2}\right)-w t\left(X_{0} \cup X_{1}\right)$, and using the fact that

$$
w t\left(X_{0} \cup X_{1}\right)=\frac{\left|X_{0}\right|}{\left|X_{0}\right|+\left|X_{1}\right|} w t\left(X_{0}\right)+\frac{\left|X_{1}\right|}{\left|X_{0}\right|+\left|X_{1}\right|} w t\left(X_{1}\right)
$$

we get that

$$
w t\left(X_{2}\right)=2 \delta_{1}+\frac{\left|X_{1}\right|}{\left|X_{0}\right|+\left|X_{1}\right|} 2 \delta_{0}+w t\left(X_{0}\right)
$$

This becomes

$$
\begin{aligned}
\epsilon & =\frac{\left|X_{2}\right|}{\left|X_{1}\right|+\left|X_{2}\right|} \delta_{1}+\left(\frac{\left|X_{2}\right|}{\left|X_{1}\right|+\left|X_{2}\right|} \frac{\left|X_{1}\right|}{\left|X_{0}\right|+\left|X_{1}\right|}+\frac{\left|X_{1}\right|}{\left|X_{1}\right|+\left|X_{2}\right|}\right) \delta_{0} \\
& \leq\left|\delta_{1}\right|+\frac{3}{2}\left|\delta_{0}\right| .
\end{aligned}
$$

Finally, we note that $\left|\delta_{0}\right|$ is $A^{\prime}$ 's advantage in distinguishing $K \backslash K^{\prime}$ from $K^{\prime}$ and $\left|\delta_{1}\right|$ is $A$ 's advantage in distinguishing $G \backslash K$ from $K^{\prime}$, which concludes the proof.

Proposition 4.3. Let $\mathcal{S M}_{(G, K)}$ and $\mathcal{S M}_{\left(G^{\prime}, K^{\prime}\right)}$ be subgroup membership problems, and let $X, X^{\prime}, Y, Y^{\prime}$ be the uniform distributions on $K, K^{\prime}, G \backslash K$ and $G^{\prime} \backslash K^{\prime}$, respectively. Suppose there is some efficient algorithm $B$ that takes elements of $G$ into elements of $G^{\prime}$, such that $B(X)$ and $X^{\prime}$ are $\epsilon$-close and $B(Y)$ and $Y^{\prime}$ are $\epsilon$-close.

Then any distinguisher $A^{\prime}$ for $\mathcal{S M}_{\left(G^{\prime}, K^{\prime}\right)}$ can be turned into a distinguisher A for $\mathcal{S M}_{(G, K)}$, and we have that

$$
\operatorname{Adv}_{A^{\prime}}^{\mathcal{S} \mathcal{M}_{\left(G^{\prime}, K^{\prime}\right)}} \leq \operatorname{Adv}_{A}^{\mathcal{S} \mathcal{M}_{(G, K)}}+\epsilon
$$

Proof. Let $A^{\prime}$ be a distinguisher for $\mathcal{S M}_{\left(G^{\prime}, K^{\prime}\right)}$. Let $A$ be the algorithm that on input of $G, K$ and $x$ simply samples $y$ from $B(x)$, runs $A^{\prime}$ on input of $G^{\prime}, K^{\prime}$ and $y$ and returns the output of $A^{\prime}$.

For brevity, we suppress the input $G, K, G^{\prime}$ and $K^{\prime}$ to $A$ and $A^{\prime}$. Note that by Lemma 2.1,

$$
\begin{aligned}
\mid \operatorname{Pr}[A(x)=1 \mid x & \leftarrow K]-\operatorname{Pr}\left[A^{\prime}(y)=1 \mid y \leftarrow K^{\prime}\right] \mid \\
& =\left|\operatorname{Pr}\left[A^{\prime}(y)=1 \mid y \leftarrow B(K)\right]-\operatorname{Pr}\left[A^{\prime}(y)=1 \mid y \leftarrow K^{\prime}\right]\right| \leq \epsilon
\end{aligned}
$$

The same holds for $G \backslash K$ and $G^{\prime} \backslash K^{\prime}$. Hence we get

$$
\begin{aligned}
\operatorname{Adv}_{A^{\prime}}^{\left.\mathcal{S} \mathcal{M}_{\left(G^{\prime}, K^{\prime}\right)}\right)} & =\frac{1}{2}\left|\operatorname{Pr}\left[A^{\prime}(y)=1 \mid y \leftarrow K^{\prime}\right]-\operatorname{Pr}\left[A(y)=1 \mid y \leftarrow G^{\prime} \backslash K^{\prime}\right]\right| \\
& \leq \frac{1}{2}\left|\operatorname{Pr}\left[A^{\prime}(x) \mid x \leftarrow K\right]-\operatorname{Pr}[A(x) \mid x \leftarrow G \backslash K]\right|+\epsilon \\
& =\operatorname{Adv}_{A}^{\mathcal{S} \mathcal{M}_{(G, K)}}+\epsilon,
\end{aligned}
$$

which concludes the proof.

### 4.2 Splitting problems

Let $\Gamma_{\tau}$ be a probability space on a set of triples $(G, K, H)$, where $G$ is a finite abelian group, and $K$ and $H$ are subgroups of $G$ such that $G=K H$ and $K \cap H=$ $\{1\}$. Let $\Gamma=\left\{\Gamma_{\tau}\right\}$.

We assume that an efficient algorithm is available for $\Gamma$ that samples $G$ (almost) uniformly at random, and we denote sampling $x$ from the resulting probability space by $x \approx G$.

For any such triple $(G, K, H) \in \Gamma_{\tau}$, any element $x \in G$ can be written uniquely as $x=x_{1} x_{2}$ with $x_{1} \in K, x_{2} \in H$. In other words, there is an isomorphism $G \xrightarrow{\sim} K \times H$. The natural question to ask is if this isomorphism is easy to compute. (The isomorphism $K \times H \xrightarrow{\sim} G$ is simply multiplication.)
Definition. The splitting problem $\mathcal{S} \mathcal{P}_{\Gamma}$ derived from the family $\Gamma$ described above is the following problem: The problem instances are tuples $(G, K, H, x)$, where $x$ is sampled uniformly at random from $G$. The answer is a pair $\left(x_{1}, x_{2}\right) \in$ $K \times H$ such that $x=x_{1} x_{2}$.

If the required sampling algorithms for $\Gamma$ are available, we can derive a subgroup membership problem $\mathcal{S M}_{\Gamma}$ from $\mathcal{S P}_{\Gamma}$, by forgetting $H$. It is obvious that $\mathcal{S} \mathcal{M}_{\Gamma}$ reduces to $\mathcal{S P}_{\Gamma}$.

Let $(G, K, H, x)$ be an instance of a splitting problem with answer $\left(x_{1}, x_{2}\right)$. We sometimes say that $\left(x_{1}, x_{2}\right)$ is the splitting of $x$ and that the projection of $x$ on $K$, respectively $H$ is $x_{1}$, respectively $x_{2}$. We also use the notation $\pi_{K}(x)$ and $\pi_{H}(x)$, for $x_{1}$ and $x_{2}$, respectively.

As for subgroup membership problems, $\mathcal{S P}_{(G, K, H)}$ denotes a subproblem of $\mathcal{S P}_{\Gamma}$.

The splitting problem has also been called projection problem [31].
Proposition 4.4. Let $\mathcal{S P}_{(G, K, H)}$ be a splitting problem. If there are efficient algorithms available to sample $K$ and $H$ (almost) uniformly at random, then the splitting problem $\mathcal{S P}_{(G, K, H)}$ is self-reducible.
Proof. Given an instance $x$, sample $z_{1}$ (almost) uniformly from $K$ and $z_{2}$ (almost) uniformly from $H$ and output the instance $x z_{1} z_{2}$ and the function $\left(y_{1}, y_{2}\right) \mapsto$ $\left(y_{1} z_{1}^{-1}, y_{2} z_{2}^{-1}\right)$. The output is (almost) uniformly distributed in $G$, and independent of the input.
Proposition 4.5. Let $\mathcal{S P}_{(G, K, H)}$ be a splitting problem such that $\operatorname{gcd}(|K|,|H|)=$ 1. Then anyone who knows $|K|$ and $|H|$ can solve the splitting problem using a deterministic and efficient algorithm.
Proof. If $m=|H|$ and $m^{\prime} \equiv|H|^{-1}(\bmod |K|)$, then $x_{1}=x^{m m^{\prime}}$, and $x_{2}=$ $x x_{1}^{-1}$.

Definition. Let $\Gamma_{\tau}$ be a probability space on a set of tuples $(G, K, H, \sigma)$, such that $\mathcal{S P}_{(G, K, H)}$ is a splitting problem and $\sigma$ is an easy to compute function that solves the splitting problem. Let $\Gamma=\left\{\Gamma_{\tau}\right\}$. We get a trapdoor splitting problem $\mathcal{T S} \mathcal{P}_{\Gamma}$. We denote the underlying splitting problem by $\mathcal{S P}_{\Gamma}$.

### 4.3 Subgroup discrete logarithm problems

For several of the splitting problems we discuss, it will be easy to compute discrete logarithms in $H$. Because of applications to cryptography, we need to ask if it is easier to compute the discrete logarithm of the projection than to compute the projection itself.

Definition. Let $\mathcal{S P}_{(G, K, H)}$ be a splitting problem, and let $g \in G$ be such that $H \subseteq\langle g\rangle$ (hence, $H$ must be cyclic). The $H$-subgroup discrete logarithm of $y \in G$ to the base $g$ is the discrete logarithm of $\pi_{H}(y)$ to the base $\pi_{H}(g)$.

Let $\Gamma_{\tau}$ be a probability space on a set of tuples $(G, K, H, g)$, such that for any tuple $(G, K, H), \mathcal{S P}_{(G, K, H)}$ is a splitting problem. (Note that this implies that there must be an algorithm available for sampling $G$ (almost) uniformly at random.) Let $\Gamma=\left\{\Gamma_{\tau}\right\}$.

Definition. The subgroup discrete logarithm problem $\mathcal{S D L}_{\Gamma}$ derived from the family $\Gamma$ described above is the following problem: The instances are tuples $(G, K, H, g, y)$, where $(G, K, H, g)$ is sampled from $\Gamma$ and $y$ is sampled uniformly from $G$. The answer is $\log _{\pi_{H}(g)} \pi_{H}(y)$.

As usual, we shall discuss subproblems $\mathcal{S D} \mathcal{L}_{(G, K, H, g)}$. We shall write the subgroup discrete logarithm of $y \in G$ as $\log _{(H, g)} y$.

The subgroup discrete logarithm problem has appeared as the partial discrete logarithm problem in the literature.

Definition. Let $\Gamma_{\tau}$ be a probability space on a set of tuples $\left(G, K, H, g, \sigma, f^{\prime}\right)$, such that $\mathcal{S D} \mathcal{L}_{(G, K, H, g)}$ is a subgroup discrete logarithm problem, $\sigma$ is an easy to compute function that solves the splitting problem, and $f^{\prime}$ is an easy to compute function that solves the discrete logarithm problem in $H$. Let $\Gamma=\left\{\Gamma_{\tau}\right\}$. We get a trapdoor subgroup discrete logarithm problem $\mathcal{T S D} \mathcal{L}_{\Gamma}$. We denote the underlying subgroup discrete logarithm problem by $\mathcal{S D} \mathcal{L}_{\Gamma}$.

It is quite obvious that if the element $g$ is in $H$, and discrete logarithms are easy to compute in $H$, then the subgroup discrete logarithm problem and the splitting problem are essentially the same.

The subgroup discrete logarithm problem is clearly very similar to the discrete logarithm problem.

It is easy to show that if one can compute subgroup discrete logarithms to some base $g$, then one can compute $|H|$. It is also easy to show that if one can compute subgroup discrete logarithms to one base, one can compute discrete logarithms to any base.

Suppose the order $|H|$ is known. If $X$ is distributed uniformly on $\{0, \ldots,|H|-$ $1\}$ and $Z$ is distributed uniformly on $K$, then $y g^{X} Z$ is distributed uniformly on $G$, it is independent of $y$ and $\log _{(H, g)} y g^{X} Z \equiv X+\log _{(H, g)} y(\bmod |H|)$. This means that the subgroup discrete logarithm problem is self-reducible when the order of $H$ is known.

We shall prove two results about the subgroup discrete logarithm problem. There are similar results for the discrete logarithm problem. Note that these reductions may or may not be useful in practice, since there is a significant increase in the work required. But since several of our bounds are quite weak, much better results could perhaps be obtained in practice, depending on the specific properties of the algorithm found.

Proposition 4.6. Let $\mathcal{S D}_{(G, K, H, g)}$ be a subgroup discrete logarithm problem, and let $H^{\prime}$ be a prime-order subgroup of $H,\left|H^{\prime}\right|=\ell$ and $\ell^{k}$ the biggest power of $\ell$ that divides $H$. Given the order $|H|$ and an oracle $A$ that distinguishes $K$ from $K H^{\prime} \backslash K$ with advantage $\epsilon$, and such that

$$
\operatorname{Pr}\left[A(x)=1 \mid x \leftarrow K H^{\prime} \backslash K\right]=\alpha,
$$

we can for any $0<\beta<1$ recover the subgroup discrete logarithm modulo $\ell^{k}$ with probability $1-k \ell \beta$ making

$$
\left\lceil\frac{4 k}{\beta \epsilon^{2}} \max \{\alpha(1-\alpha),(\alpha+\epsilon)(1-\alpha-\epsilon)\}\right\rceil
$$

queries to $A$.
Proof. Let $H^{\prime}$ be a subgroup of $H$ of prime order $\ell$, and let $h=g^{|H| / \ell \mid}$. By Proposition 4.1, $\mathcal{S} \mathcal{M}_{\left(K H^{\prime}, K\right)}$ is self-reducible, so we may assume that $\operatorname{Pr}[A(x)=$ $\left.1 \mid \log _{\left(H^{\prime}, h\right)}=i\right]$ is independent of $i$ when $i=1, \ldots, \ell-1$, and that $\operatorname{Pr}[A(x)=$ $1 \mid x \leftarrow K]=\operatorname{Pr}\left[A(x)=1 \mid x \leftarrow K H^{\prime} \backslash K\right]+\epsilon$.

Let $x$ be an instance of the subgroup discrete logarithm problem. Raising $x$ to the $|H| / \ell$ th power will turn $x$ into an instance $x^{\prime}$ of the subgroup discrete logarithm problem $\mathcal{S D} \mathcal{L}_{\left(K H^{\prime}, K, H^{\prime}, h\right)}$. If the subgroup discrete logarithm of $x^{\prime}$ is $m$, then $x^{\prime} h^{-m} \in K$. This means that we can use $A$ to test if $\log _{\left(H^{\prime}, h\right)} x \equiv m$ $\left(\bmod \left|H^{\prime}\right|\right)$.

Each test is a Bernoulli trial, so let $\alpha=\operatorname{Pr}\left[A(x)=1 \mid x \in K H^{\prime} \backslash K\right]$. We repeating each trial $s$ times and let $X$ be the number of times 1 occurs. When we test the wrong $m$, we expect the number of ones to be $s \alpha$, while when we test the correct $m$, we expect $s(\alpha+\epsilon)$ ones.

When we test the wrong $m$, we get that

$$
\operatorname{Pr}[|X / s-\alpha|<\epsilon / 2] \geq 1-\frac{\alpha(1-\alpha)}{s \epsilon^{2} / 4} .
$$

If we fix some maximal failure probability $\beta$, we get that

$$
s_{0}=\frac{\alpha(1-\alpha)}{\beta \epsilon^{2} / 4} .
$$

This means that, except with probability $\beta$, after so many Bernoulli trials, the proportion of ones will be within $\epsilon / 2$ of $\alpha$.

When we test the correct $m$, we get that

$$
\operatorname{Pr}[|X / s-\alpha-\epsilon|<\epsilon / 2] \geq 1-\frac{(\alpha+\epsilon)(1-\alpha-\epsilon)}{s \epsilon^{2} / 4}
$$

Again, if we fix some maximal failure probability $\beta$, we get that

$$
s_{1}=\frac{(\alpha+\epsilon)(1-\alpha-\epsilon)}{\beta \epsilon^{2} / 4} .
$$

So we make a total of $s=\max \left\{s_{1}, s_{2}\right\}$ tests for each possible value of the subgroup discrete logarithm problem, and choose the value that gives the most ones.

The total probability that this test is successful is at least $(1-\beta)^{\ell-1}(1-\beta) \geq$ $1-\ell \beta$.

To conclude the proof, we note that an $\ell$-adic lift as used in the PohligHellman algorithm can be used to recover the discrete logarithm modulo the prime power.

Suppose $|H|$ is smooth and square-free. Let $H^{\prime}$ be a prime-ordered subgroup of $H$, and let $H^{\prime \prime}$ be the subgroup of $H$ of order $|H| /\left|H^{\prime}\right|$. It is quite clear that Proposition 4.3 applies. This means that if we can argue that all of the subgroups of $H$ are equal, in the sense that either the subgroup membership problem $\mathcal{S M}_{\left(G, K H^{\prime \prime}\right)}$ is difficult for all prime-order subgroups $H^{\prime}$ of $H$, or for none, the subgroup discrete logarithm problem reduces to the subgroup membership problem $\mathcal{S M}_{(G, K)}$.

Let $\mathcal{S D} \mathcal{L}_{(G, K, H, g)}$ be a subgroup discrete logarithm problem. Let $d$ be an integer relatively prime to $|H|$, and define the function

$$
l s d(x)=\left(\log _{H, g} x\right) \bmod d
$$

It is quite clear that $l s d$ has a limited self-reducibility property, in that for any $x \in G$ and $x^{\prime} \in K, l s d(x)=l s d\left(x x^{\prime}\right)$. But the success probability of an algorithm for computing $l s d$ may depend on the value of $\log _{H, g} x$, so we need to be a bit more careful.

Let $\mathcal{P}=(P, X, S, f)$ be a problem, and suppose that the set of answers $S$ has a group structure (written additively). Let $A$ be an algorithm that tries to solve $\mathcal{P}$. We say that $A$ has systematic advantage

$$
\begin{aligned}
\operatorname{SAdv}_{A}^{\mathcal{P}}= & \operatorname{Pr}[A(x)=\mathcal{P}(x) \mid x \leftarrow \mathcal{P}]- \\
& \max _{s \in S \backslash\{0\}} \operatorname{Pr}[A(x)=\mathcal{P}(x)+s \mid x \leftarrow \mathcal{P}] .
\end{aligned}
$$

Let $d$ be fixed. For a subgroup discrete logarithm problem $\mathcal{S D} \mathcal{L}_{(G, K, H, g)}$, we get the problem of computing $l s d$. For $\mathcal{S D} \mathcal{L}_{\Gamma}$, we define the systematic advantage to be

$$
\operatorname{SAdv}_{A}^{\mathcal{S D} \mathcal{L}_{\Gamma}}(\tau)=\mathrm{E}\left[\operatorname{SAdv}_{A}^{\mathcal{S D} \mathcal{L}_{(G, K, H, g)}} \mid(G, K, H, g) \leftarrow \Gamma_{\tau}\right] .
$$

Note that if $d=2$, then systematic advantage reduces to advantage as defined in (2.1).

We shall prove the following result.
Proposition 4.7. Let $\mathcal{S D} \mathcal{L}_{\Gamma}$ be a subgroup discrete logarithm problem, and let $d$ be a positive integer such that for every subproblem $\mathcal{S D} \mathcal{L}_{(G, K, H, g)}$ of $\mathcal{S D} \mathcal{L}_{\Gamma}, d$ is relatively prime to $|H|$ and it is easy to find some number e such that de $\equiv 1$ $(\bmod |H|)$. If there is some efficient algorithm for computing lsd with significant systematic advantage, then $\mathcal{S D} \mathcal{L}_{\Gamma}$ is not hard.

The case $d=2$ was proved in [8]. We have a somewhat more general result, so we include the proofs here. We first prove two lemmas. The first says that the systematic advantage gives us a stronger property for a restricted problem.

We remark that even if we do not know the order $|H|$, but we know an integer $n$ such that $\operatorname{gcd}(n, d)=1$ and $|H|$ divides $n$, then we can find $e$ by computing inverses modulo $n$ instead of $|H|$.
Lemma 4.8. Let $\mathcal{S D} \mathcal{L}_{(G, K, H, g)}$ be a subgroup discrete logarithm problem, and let $A$ be an algorithm that computes lsd with systematic advantage $\epsilon>0$.

Then there is an algorithm $A^{\prime}$ that computes lsd such that for all $x \in G$ with $\log _{H, g} x \leq\lfloor|H| \epsilon / 3\rfloor$,

$$
\epsilon / 3 \leq \operatorname{Pr}\left[T_{0}(x)\right]-\max _{0<s<d} \operatorname{Pr}\left[T_{s}(x)\right],
$$

where $T_{s}(x)$ is the event that $A^{\prime}(x) \equiv l s d(x)+s(\bmod d)$.
Proof. The algorithm $A^{\prime}$ takes as input $x$, samples $X$ uniformly from $K$ and $Y$ uniformly from $\{0, \ldots,|H|-1\}$ and outputs $\left(A\left(X x g^{Y}\right)-Y\right) \bmod d$.

Let $X$ be uniformly distributed in $K, Y$ be uniformly distributed in $\{0, \ldots$, $|H|-1\}$, and let $x \in G$ be such that $\log _{H, g} x \leq\lfloor|H| \epsilon / 3\rfloor$. Then $\log _{H, g} X x g^{Y}=$ $Y+\log _{H, g} x$ with probability at least $1-\epsilon / 3$, and $\log _{H, g} X x g^{Y}=Y+\log _{H, g} x-|H|$ with probability at most $\epsilon / 3$.

Let $Z=X x g^{Y}$, and set $\alpha=\max _{0<s<d} \operatorname{Pr}\left[T_{s}(x) \mid x \leftarrow G\right]$. We compute

$$
\begin{aligned}
\operatorname{Pr}\left[T_{0}(Z)\right]= & \operatorname{Pr}\left[T_{0}(Z) \mid Y \leq\lfloor|H| \epsilon / 3\rfloor\right](1-\epsilon / 3) \\
& +\operatorname{Pr}\left[T_{0}(Z) \mid Y>\lfloor|H| \epsilon / 3\rfloor\right] \epsilon / 3 \\
\geq & \operatorname{Pr}\left[T_{0}(Z) \mid Y \leq\lfloor|H| \epsilon / 3](1-\epsilon / 3)\right. \\
\geq & \frac{\alpha+\epsilon-\epsilon / 3}{1-\epsilon / 3}(1-\epsilon / 3)=\alpha+2 \epsilon / 3,
\end{aligned}
$$

and

$$
\begin{aligned}
\operatorname{Pr}\left[T_{s}(Z)\right]= & \operatorname{Pr}\left[T_{s}(Z) \mid Y \leq\lfloor|H| \epsilon / 3\rfloor\right](1-\epsilon / 3) \\
& +\operatorname{Pr}\left[T_{s}(Z) \mid Y>\lfloor|H| \epsilon / 3\rfloor\right] \epsilon / 3 \\
\leq & \frac{\alpha}{1-\epsilon / 3}(1-\epsilon / 3)+\epsilon / 3=\alpha+\epsilon / 3 .
\end{aligned}
$$

The result follows.

Now we show that a simple majority vote-type algorithm can amplify the success probability of an oracle for $l s d$.

Lemma 4.9. Let $A^{\prime}$ be the algorithm in Lemma 4.8. Then there is an algorithm $A^{\prime \prime}$ that computes lsd with success probability $\beta$ using less than $O\left(1 /\left(\epsilon^{8}(1-\beta)\right)\right)$ (or $O\left(1 /\left(\epsilon^{3}(1-\beta)\right)\right)$ if $d$ is small) invocations of $A^{\prime}$.

Proof. The algorithm $A^{\prime \prime}$ takes as input $x$ and runs $A^{\prime}$ several times, counting the different answers. Then it returns the most frequent answer.

First of all, we may suppose that 0 is the correct answer. Let $X_{i}, 0 \leq i<d$, count the number of times $A^{\prime}$ returns the value $i$. Since different runs of $A^{\prime}$ are independent, $\left(X_{0}, X_{1}, \ldots, X_{d-1}\right)$ will be a multinomially distributed random variable with parameters $\left(\alpha_{0}, \alpha_{1}, \ldots, \alpha_{d-1}\right)$.

By assumption, $\alpha_{0} \geq \alpha_{i}+\epsilon / 3$ for $1 \leq i<d$. If $d$ is smaller than $O(1 / \epsilon)$, the result follows immediately from Lemma 2.6.

So suppose $d>4 / \epsilon$. Let $\left\{S_{1}, S_{2}, \ldots, S_{d^{\prime}}\right\}$ be a partition of $\{1,2, \ldots, d-1\}$, and let $Z_{i}=\sum_{s \in S_{i}} X_{s}$ and $\alpha_{i}^{\prime}=\sum_{s \in S_{i}} \alpha_{s}$. Let $Z_{0}=X_{0}$ and $\alpha_{0}^{\prime}=\alpha_{0}$. It is possible to find a partition such that there is at most one $\alpha_{i}^{\prime}<\epsilon / 4$ and

$$
\alpha_{0}^{\prime} \geq \alpha_{i}^{\prime}+\epsilon / 2, \quad i=1, \ldots, d^{\prime}-1
$$

It is clear that $\left(Z_{0}, \ldots, Z_{d^{\prime}-1}\right)$ is a multinomially distributed random variable with parameters $\left(\alpha_{0}^{\prime}, \ldots, \alpha_{d^{\prime}-1}^{\prime}\right)$, and $d^{\prime} \leq 4 / \epsilon$. It is also clear that

$$
\operatorname{Pr}\left[X_{0}=\max \left\{X_{0}, \ldots, X_{d-1}\right\}\right] \geq \operatorname{Pr}\left[Z_{0}=\max \left\{Z_{0}, \ldots, Z_{d^{\prime}-1}\right\}\right],
$$

and the result again follows from Lemma 2.6.
Proof of Proposition 4.7. Note that if we know $l s d(x)$, then $\log _{H, g} x g^{-l s d(x)}$ is divisible by $d$. We can divide by $d$ by computing the $e$ th power. To recover every $d$-adic digit of $\log _{H, g} x$ in this way requires $\left\lceil\log _{d}|H|\right\rceil$ computations of $l s d(x)$.

If every computation of $l s d(x)$ fails with probability $\gamma$, then the entire computation is correct with probability $(1-\gamma)^{\left\lceil\log _{d}|H|\right\rceil} \approx 1-\left\lceil\log _{d}|H|\right\rceil \gamma$, so the maximal failure rate of the $l s d$ computation is polynomial in $\log _{d}|H|$.

Lemma 4.8 and Lemma 4.9 say that we can compute $l s d$ with sufficiently small failure rate, keeping the number of invocations of $A$ polynomial in $1 / \epsilon$ and $\log _{d}|H|$, if the subgroup discrete logarithm is restricted.

Finally, since the probability that an instance sampled uniformly satisfies this restriction is polynomial in $1 / \epsilon$, the self-reducibility property of the subgroup discrete logarithm problem allows us to lift the restriction. This concludes the proof.

Given an integer $d$ relatively prime to $|H|$, we can write

$$
\log _{H, g} x=\sum_{i} \alpha_{i} d^{i}
$$

We can define $l s d_{i}(x)=\alpha_{i}$.

Proposition 4.10. Let $\mathcal{S D}_{(G, K, H, g)}$ be a subgroup discrete logarithm problem. If there is some efficient algorithm A for computing lsd $i_{i}$ with significant systematic advantage, then there is an algorithm for computing subgroup discrete logarithms when the answers are restricted to be less than $d^{\left\lfloor\log _{d}|H|\right\rfloor-i}$, using $A$ a polynomial (in $\log |H|$ ) number of times.
Proof. If we know that $\log _{H, g} x$ is less than $d^{l}$, then for any $i<\left\lfloor\log _{d}|H|\right\rfloor-l$, we have that $\log _{H, g} x^{d^{i}}=d^{i} \log _{H, g} x$, so $l s d(x)=l s d_{i}\left(x^{d^{i}}\right)$. In order to apply Lemma 4.9 and proceed as in the proof of Proposition 4.7, we need to show how to randomise the queries to the algorithm $A$.

Following Lemma 4.8, we assume that $\log _{H, g} x \leq\lfloor|H| \epsilon / 3\rfloor$. Then we sample $r$ uniformly at random from $\left\{0, \ldots, d^{i}-1\right\}$, and $r^{\prime}$ uniformly at random from $\left\{d^{i}, \ldots,|H|-1\right\}$. It is clear that with probability $1-\epsilon / 3, d^{i} \log _{H, g} x+r+r^{\prime}<|H|$, and that $A$ has systematic advantage at least $\epsilon / 3$ on the reduced instance.

The interesting case is for $d=2$. From the above proof, it is clear that we can give the algorithm $A$ that tries to compute the $i$ th bit of the subgroup discrete logarithm the lower order bits, since except with probability $\epsilon / 3$, they are exactly $r$. This means that all of the bits are simultaneously hard to compute.

We can consider this informally as a limiting process. As the restriction on the subgroup discrete logarithm decreases, the number of simultaneously hard bits increases. In the limit, where it is difficult to prove that the subgroup discrete logarithm is zero, given the knowledge that it is zero, we get the subgroup membership problem.

### 4.4 Symmetric subgroup membership problems

Let $\Gamma$ be defined as in Section 4.2, and suppose that there is an efficient algorithm available that samples $K$ and $H$ (almost) uniformly at random. As usual, sampling using this algorithm is denoted by $x \approx K$ and $x \approx H$. It is clear that we can derive two families $\Gamma^{\prime}$ and $\Gamma^{\prime \prime}$ from $\Gamma$ by forgetting in turn $H$ and $K$, and that we can derive subgroup membership problems from these families.
Definition. The symmetric subgroup membership problem $\mathcal{S S} \mathcal{M}_{\Gamma}$ is the double problem $\mathcal{S} \mathcal{M}_{\Gamma^{\prime}}$ and $\mathcal{S} \mathcal{M}_{\Gamma^{\prime \prime}}$. We say that the $\mathcal{S} \mathcal{M}_{\Gamma}$ is hard if both $\mathcal{S} \mathcal{M}_{\Gamma^{\prime}}$ and $\mathcal{S M}_{\Gamma^{\prime \prime}}$ are hard.

Note that $\mathcal{S S} \mathcal{M}_{\Gamma}$ is not a problem in the sense defined in Section 2.2, but it is convenient to borrow the terminology. As for subgroup membership problem, we usually consider subproblems $\mathcal{S S}_{(G, K, H)}$.
Theorem 4.11. Let $\mathcal{S S}_{(G, K, H)}$ be a symmetric subgroup membership problem such that $G$ is cyclic, and let $A$ be an algorithm that decides the Decision DiffieHellman problem on $G$. Then for any $\delta^{\prime}>0$ and certain algorithms $A_{1}, A_{2}$, $A_{3}, A_{1}^{\prime}, A_{2}^{\prime}$ and $A_{3}^{\prime}$ that use $A$ once as an oracle and otherwise do $O\left(\log 1 / \delta^{\prime}\right)$ exponentiations in $G$, we have that

$$
\operatorname{Adv}_{A}^{\mathcal{D D H}_{G}} \leq \operatorname{Adv}_{A_{1}}^{\mathcal{S M}_{(G, K)}}+\operatorname{Adv}_{A_{2}}^{\mathcal{S M}_{(G, K)}}+\operatorname{Adv}_{A_{3}}^{\mathcal{S M}_{(G, H)}}+f(|K|)+O\left(\delta+\delta^{\prime}\right)
$$

and

$$
\operatorname{Adv}_{A}^{\mathcal{D D} \mathcal{H}_{G}} \leq \operatorname{Adv}_{A_{1}^{\prime}}^{\mathcal{S} \mathcal{M}_{(G, H)}}+\operatorname{Adv}_{A_{2}^{\prime}}^{\mathcal{S} \mathcal{M}_{(G, H)}}+\operatorname{Adv}_{A_{3}^{\prime}}^{\mathcal{S} \mathcal{M}_{(G, K)}}+f(|H|)+O\left(\delta+\delta^{\prime}\right)
$$

where the sampling algorithms for $K$ and $H$ are $\delta$-close to uniform, and

$$
f(x)=\frac{|G|-\phi(|G|)}{|G|}+\frac{4 x}{|G|}+\frac{x-\phi(x)}{x} .
$$

We shall need the following result quite often, so we state it as a separate lemma.

Lemma 4.12. Let $G$ be a finite cyclic group, and let $K$ and $H$ be non-trivial subgroup of $G$ such that $K \cap H=\{1\}$ and $G=K H$. Let $g$ be a generator for $K$. Let $X$ be uniformly distributed on $\{0, \ldots,|G|-1\}, Y$ be uniformly distributed on $G$, while $Z$ is uniformly distributed on $H$. Let $U=\left(g^{X}, Y^{X}\right)$ and let $V=$ $\left(g^{X}, Y^{X} Z\right)$. Then

$$
\operatorname{Dist}(U, V) \leq 1-\frac{\phi(|H|)}{|H|}
$$

Proof. Let $X_{1}=X \bmod |K|$ and $X_{2}=X \bmod |H|$, and let $Y=Y_{1} Y_{2}$, where $Y_{1} \in K$ and $Y_{2} \in H$. It is clear that $X_{1}, X_{2}, Y_{1}$ and $Y_{2}$ are independently and uniformly distributed, and that

$$
U=\left(g^{X_{1}}, Y_{1}^{X_{1}} Y_{2}^{X_{2}}\right) \text { and } V=\left(g^{X_{1}}, Y_{1}^{X_{1}} Y_{2}^{X_{2}} Z\right)
$$

If we let $U^{\prime}=Y_{2}^{X_{2}}$ and $V^{\prime}=Y_{2}^{X_{2}} Z$, it is clear that

$$
\operatorname{Dist}(U, V)=\operatorname{Dist}\left(U^{\prime}, V^{\prime}\right)
$$

and that $V^{\prime}$ is uniformly distributed on $H$.
Now we condition on a fixed $y_{2} \in H$. If this $y_{2}$ is a generator, then it is clear that $y_{2}^{X_{2}}$ is uniformly distributed on $H$. By Lemma 2.5 , it follows that

$$
\operatorname{Dist}\left(U^{\prime}, V^{\prime}\right) \leq 1-\frac{\phi(|H|)}{|H|}
$$

which concludes the proof.
Proof of Theorem 4.11. $A$ is an algorithm that takes as input $G, c_{1}, c_{2}, c_{3}, c_{4}$, and returns 0 or 1. We shall need the experiments listed in Figure 4.1.

The idea is to make small changes to the input distribution of $A$, interleaved with changes that correspond to the subgroup membership problems. Then we shall show that in the end, $A$ has no advantage against the final input distribution, and the only significant changes in advantage correspond to the subgroup membership problems. Therefore, if $A$ has significant advantage, at least one of the derived algorithms must have significant advantage against a subgroup membership problem.

## Experiment 1.

Input: $A, G$.

1. $u, v, w \leftarrow\{0, \ldots,|G|-1\}$.
2. $x \leftarrow G$.
3. $b \leftarrow\{0,1\}$.
4. If $b=1$, then $z \leftarrow x^{u v}$, otherwise $z \leftarrow x^{w}$.
5. $b^{\prime} \leftarrow A\left(x, x^{u}, x^{v}, z\right)$.
6. If $b=b^{\prime}$, then output 1 , otherwise output 0 .

## Experiment 2.

Input: $A, G, x \in G$.

1. $u, v, w \leftarrow\{0, \ldots,|G|-1\}$.
2. $b \leftarrow\{0,1\}$.
3. If $b=1$, then $z \leftarrow x^{u v}$, otherwise $z \leftarrow x^{w}$.
4. $b^{\prime} \leftarrow A\left(x, x^{u}, x^{v}, z\right)$.
5. If $b=b^{\prime}$, then output 1 , otherwise output 0 .

Output: 0 or 1.
Output: 0 or 1.

## Experiment 3.

Input: $A, G, y \in G$.

1. $u, v, w \leftarrow\{0, \ldots,|G|-1\}$.

## Experiment 4.

Input: $A, G, h \in G$.

1. $u, v, w \leftarrow\{0, \ldots,|G|-1\}$.
2. $x \leftarrow K$.
3. $x \leftarrow K, y \leftarrow G$.
4. $b \leftarrow\{0,1\}$.
5. If $b=1$, then $z \leftarrow y^{v}$, otherwise $z \leftarrow y^{w}$.
6. $b^{\prime} \leftarrow A\left(x, x^{v}, y, z\right)$.
7. If $b=b^{\prime}$, then output 1 , otherwise output 0 .

Output: 0 or 1.
3. $b \leftarrow\{0,1\}$.
4. If $b=1$, then $z \leftarrow h y^{v}$, otherwise $z \leftarrow y^{w}$.
5. $b^{\prime} \leftarrow A\left(x, x^{v}, y, z\right)$.
6. If $b=b^{\prime}$, then output 1 , otherwise output 0 .

Output: 0 or 1.

Figure 4.1: Experiments for the proof of Theorem 4.11.

First of all, we note that by using the sampling algorithms of the symmetric subgroup membership problem, and by Lemma 2.3, even if we do not know $|G|$, we can implement modified experiments whose underlying probability spaces are $O\left(\delta+\delta^{\prime}\right)$-close to the original experiment, for any $\delta^{\prime}>0$. The extra work needed is $O\left(\log 1 / \delta^{\prime}\right)$ exponentiations in $G$.

Consider first Experiment 1. Let $T_{1}$ denote the event that the experiment returns 1. If the $x$ chosen at random is a generator, then it is clear that the experiment runs exactly as Experiment 2 in Section 2.3. By Lemma 2.5, we get that

$$
\begin{equation*}
\operatorname{Adv}_{A}^{\mathcal{D} \mathcal{D H}_{G}} \leq\left|\operatorname{Pr}\left[T_{1}\right]-1 / 2\right|+\frac{|G|-\phi(|G|)}{|G|} . \tag{4.1}
\end{equation*}
$$

Next we consider Experiment 2. Suppose the input $x$ to Experiment 2 is sampled uniformly from $G \backslash K$. Let $T_{2}$ be the event that the experiment returns 1 in this case. The only difference now between Experiment 1 and Experiment 2 is that Experiment 1 samples $x$ uniformly from $G$, not $G \backslash K$. By Lemma 2.2, we see that

$$
\begin{equation*}
\left|\operatorname{Pr}\left[T_{1}\right]-\operatorname{Pr}\left[T_{2}\right]\right| \leq \frac{2|K|}{|G|} \tag{4.2}
\end{equation*}
$$

Now consider the case when the input $x$ to Experiment 2 is sampled uniformly from $K$, and let $T_{2}^{\prime}$ be the event that the experiment returns 1 in this case. It is clear that we can construct an algorithm $A_{1}$ for deciding $\mathcal{S M}_{(G, K)}$ from Experiment 2 such that

$$
\begin{equation*}
\operatorname{Adv}_{A_{1}}^{\mathcal{S} \mathcal{M}_{(G, K)}}=\left|\operatorname{Pr}\left[T_{2}\right]-\operatorname{Pr}\left[T_{2}^{\prime}\right]\right| . \tag{4.3}
\end{equation*}
$$

Next, we consider Experiment 3. Suppose the input $y$ to Experiment 3 is sampled uniformly from $K$. Let $T_{3}^{\prime}$ be the event that the experiment returns 1 in this case. It is clear that if the $x$ sampled is a generator for $K$, then Experiment 2 and Experiment 3 proceed identically. So

$$
\begin{equation*}
\left|\operatorname{Pr}\left[T_{3}^{\prime}\right]-\operatorname{Pr}\left[T_{2}^{\prime}\right]\right| \leq \frac{|K|-\phi(|K|)}{|K|} \tag{4.4}
\end{equation*}
$$

Let $T_{3}$ be the event that the experiment returns 1 when the input $y$ to Experiment 3 is sampled uniformly from $G \backslash K$. Again, it is clear that we can construct an algorithm $A_{2}$ for deciding $\mathcal{S M}_{(G, K)}$ from Experiment 3 such that

$$
\begin{equation*}
\operatorname{Adv}_{A_{2}}^{\mathcal{S} \mathcal{M}_{(G, K)}}=\left|\operatorname{Pr}\left[T_{3}\right]-\operatorname{Pr}\left[T_{3}^{\prime}\right]\right| \tag{4.5}
\end{equation*}
$$

Finally, we consider Experiment 4. Suppose the input $h$ to Experiment 4 is sampled uniformly from $H$. Let $T_{4}$ be the event that Experiment 4 returns 1 in this case. The only difference between $T_{4}$ and $T_{3}$ is that in Experiment 3, $y$ is sampled uniformly from $G \backslash K$, while in Experiment 4 it is sampled uniformly from $G$, and the final result is multiplied by $h$. By Lemma 2.2 and Lemma 4.12, we have that

$$
\begin{equation*}
\left|\operatorname{Pr}\left[T_{3}\right]-\operatorname{Pr}\left[T_{4}\right]\right| \leq \frac{2|K|}{|G|}+\frac{|H|-\phi(|H|)}{|H|} . \tag{4.6}
\end{equation*}
$$

Again, we can construct an algorithm $A_{3}$ for deciding $\mathcal{S} \mathcal{M}_{(G, H)}$ from Experiment 4 such that

$$
\begin{equation*}
\operatorname{Adv}_{A_{3}}^{\mathcal{S} \mathcal{M}_{(G, H)}}=\left|\operatorname{Pr}\left[T_{4}\right]-\operatorname{Pr}\left[T_{4}^{\prime}\right]\right| . \tag{4.7}
\end{equation*}
$$

If the input $h$ to Experiment 4 is sampled uniformly from $G \backslash H$, then it is clear that the input $z$ to $A$ in Experiment 4 is uniformly distributed in $G$, no matter what the value of $b$ is. So if $T_{4}^{\prime}$ is the event that Experiment 4 returns 1 when $h$ is sampled uniformly from $G \backslash H$, we know that

$$
\begin{equation*}
\operatorname{Pr}\left[T_{4}^{\prime}\right]=\frac{1}{2} \tag{4.8}
\end{equation*}
$$

Combining equations (4.1)-(4.8) proves the first claim, and the second claim follows by symmetry.

### 4.5 A catalogue

We shall discuss several subgroup membership problems that have appeared in the literature, and some new variations on those problems.

All of these subgroup membership problems have corresponding splitting problems. Some of them are interesting subgroup discrete logarithm problems, and some are possibly hard symmetric subgroup membership problems. The splitting problems and subgroup discrete logarithm problems can also be turned into trapdoor problems, and in general Proposition 4.5 allows for efficient splitting.

For every problem, we also discuss how we can sample elements from the relevant subgroups/subsets. We also discuss briefly the relationship with other problems.

### 4.5.1 Quadratic Residue problem

Let $p$ and $q$ be primes congruent to 3 modulo 4 , and let $n=p q$. Let

$$
J_{n}=\left\{x \in \mathbb{Z}_{n}^{*} \left\lvert\,\left(\frac{x}{n}\right)=1\right.\right\} .
$$

Let $Q_{n}$ be the subgroup of quadratic residues of $J_{n}$. Then we have a subgroup membership problem $\mathcal{S M}_{\left(J_{n}, Q_{n}\right)}$ and a splitting problem $\mathcal{S P}_{\left(J_{n}, Q_{n},\langle-1\rangle\right)}$. This is the Quadratic Residue problem $\mathcal{Q R}_{n}$, which first appeared in [22].

Since $\langle-1\rangle$ has order 2 , the subgroup membership problem and the splitting problem are equivalent.

While anyone who knows the order of $Q_{n}$ can solve the splitting problem, it is quicker to compute the Legendre symbol modulo either $p$ or $q$.

We can sample elements uniformly from $K$ by sampling elements uniformly from $\mathbb{Z}_{n}^{*}$ and squaring them. $H$ consists of 1 and -1 , so we can sample uniformly from
$H$ too. Note that if one can sample uniformly and independently from $K$ and $H$, then one can sample uniformly from both $G$ and $G \backslash K$.

If is obvious that anyone who can factor the modulus $n$ can solve $\mathcal{Q} \mathcal{R}_{n}$. It is not known if these are equivalent, but the best known general attack against $\mathcal{Q R}_{n}$ is to factor the modulus.

### 4.5.2 Higher Residue problem

Let $p$ and $q$ be primes congruent to 3 modulo 4 , and let $n=p q$. Suppose $c$ is an odd prime such that $c \mid p-1$, and $c \nmid q-1$. Then, in an analogue of the quadratic residue problem, we have the Higher Residue problem, where $H$ is a subgroup of $Q_{n}$ of order $c$. This was first investigated by Cohen Benaloh for non-prime but smooth $c$ ([6] and the references therein). We get a subgroup membership problem and a splitting problem. If $c$ is small, we also get a natural subgroup discrete logarithm problem.

The natural generalisation is to let $p=2 a c+1, q=2 b d+1$ and $n=p q$, $a, b$ prime, such that $\operatorname{gcd}(a c, b d)=\operatorname{gcd}(a b, c d)=1$. This was investigated by Naccache and Stern [29]. We let $G=Q_{n}, K$ be the unique subgroup of order $a b$ and $H$ be the subgroup of order $c d$. To get a natural subgroup discrete logarithm problem, we let $c$ and $d$ be smooth integers. Discrete logarithms in $H$ can then be computed using the Pohlig-Hellman algorithm. It is natural to assume that $c d$ is known. We denote the subgroup membership problem $\mathcal{S} \mathcal{M}_{(G, K)}$ by $\mathcal{H} \mathcal{R}_{n}$.

To sample elements of $K$, we can sample elements uniformly at random from $\mathbb{Z}_{n}^{*}$, and then apply the map $x \mapsto x^{2 c d}$. Note that this map is an automorphism on $K$, but on $H$ it is the zero homomorphism.

Suppose we are given an element $x \in H$. Knowing the integer $c d$ and its prime factorisation, finding the order $k$ of $x$ is easy. For any prime $\ell$ dividing $k$ and any representative $m$ for the residue class $x$, it is clear that $m^{k / \ell}$ is congruent to 1 modulo either $p$ or $q$, but not both, so $\operatorname{gcd}\left(m^{k / \ell}-1, n\right)$ gives us the factorisation of $n$.

Therefore, we cannot give away elements of $H$. In order to sample elements of $G \backslash K$, we publish a generator $g$ for $G$. By sampling $u$ uniformly from $\{1, \ldots, c d-$ $1\}$ and $r$ uniformly from $K, r g^{u}$ is uniformly distributed in $G \backslash K$.

Under the assumption that $c$ and $d$ are known (not only $c d$ ), Naccache and Stern [29] proposed an attack that required approximately

$$
\begin{equation*}
O\left(n^{1 / 2} /(c d)^{2}\right) \tag{4.9}
\end{equation*}
$$

work. This means that if $c d$ is too large, there are better algorithms for factoring the modulus than the general algorithms.

The basic idea of the attack is as follows. We have that $n=4 a b c d+2(a c+$ $b d)+1$. Then we get that

$$
n-1 \equiv 2 a c \quad(\bmod d) \text { and } n-1 \equiv 2 b d \quad(\bmod c) .
$$

In other words, knowledge of $c$ and $d$ gives us some knowledge about $a$ and $b$ that speeds up the search. If $c$ and $d$ are sufficiently small, factoring $n$ seems to be the best attack on the subgroup membership problem $\mathcal{S M}_{(G, K)}$.

Having a generator for $H$ is sometimes desirable. Instead of having $c$ and $d$ be relatively prime, as above, we can let $c=d$, so that $p=2 a c+1, q=2 b c+1$, and $n=4 a b c^{2}+2 c(a+b)+1$. Note that $c \mid(n-1)$, so we may as well assume that $c$ is known.

The interesting thing about the subgroup of order $c^{2}$ is that it is non-cyclic. If an element $[w]$ of $H$ is made public, and $w \neq 1(\bmod p)$ and $w \neq 1(\bmod q)$, then no matter what group operations are performed on $[w]$, the result will never yield a factor of $n$. It therefore seems safe to make elements of order $c$ public.

Every element of order $c$ generates a subgroup. If we have two elements $x$ and $y$ that generate distinct subgroups, we can use a Pollard $\rho$-type attack to generate a collision modulo $p$ and not modulo $q$. This means that we should not publish elements from more than one subgroup.

In general, we will be interested in the modified group structure where $K$ is the subgroup of order $a b, H$ is a cyclic subgroup of order $c$ generated by an element $g$, and $G=K H$. We denote the subgroup membership problem $\mathcal{S M}_{(G, K)}$ by $\mathcal{H} \mathcal{R}_{n}^{\prime}$.

As above, we can sample uniformly at random from $K$ by sampling uniformly at random from $\mathbb{Z}_{n}^{*}$, then applying the map $x \mapsto x^{2 c}$. We publish a generator $g$ for $H$, so sampling uniformly at random from $H$ is possible.

If we can recover the product $a b$, then we can factor $n$ given any element $m$ with Jacobi symbol -1 modulo $n$, because $m^{a b c}$ is congruent to 1 modulo one of the prime factors, and -1 modulo the other. Therefore, $\operatorname{gcd}\left(m^{a b c}-1, n\right)$ gives the factorisation of $n$.

First of all, $n-1 \equiv 0(\bmod c)$, so the attack of Naccache and Stern described above does not work. We get that

$$
\frac{n-1}{4 c^{2}}=a b+\frac{a+b}{2 c}
$$

This means that $a b$ is relatively close to $(n-1) /\left(4 c^{2}\right)$. Note that $a c \approx b c \approx \sqrt{n}$. We can use a Baby-step Giant-step type algorithm to recover $a b$ in

$$
O(\sqrt{(a+b) / c})=O\left(\sqrt{\sqrt{n} / c^{2}}\right)=O\left(n^{1 / 4} / c\right)
$$

steps. This is not bad compared with (4.9).
There are obvious elliptic curve analogues of the above group structures, based on elliptic curves modulo composite integers. Let $n$ be a product of two primes $p$ and $q$. Find elliptic curves $E^{\prime}$ and $E^{\prime \prime}$ defined over $\mathbb{F}_{p}$ and $\mathbb{F}_{q}$, respectively, such that $\# E\left(\mathbb{F}_{p}\right)=a c$ and $\# E\left(\mathbb{F}_{q}\right)=b d$. Then we use the Chinese remainder theorem to derive an elliptic curve $E$ defined over $\mathbb{Z}_{n}$ such that $\# E\left(\mathbb{Z}_{n}\right)=a b c d$.

Hasse's theorem says that $\# E\left(\mathbb{F}_{p}\right)=p+1-t_{p}\left(t_{p}\right.$ is the trace of the Frobenius endomorphism), and that $-2 \sqrt{p} \leq t_{p} \leq 2 \sqrt{p}$ (the same holds for $q$ ). The trace is distributed roughly uniformly in this range. We get that

$$
\begin{aligned}
\# E\left(\mathbb{Z}_{n}\right) & =\# E^{\prime}\left(\mathbb{F}_{p}\right) \# E^{\prime \prime}\left(\mathbb{F}_{q}\right)=a b c d \\
& =\left(p+1-t_{p}\right)\left(q+1-t_{q}\right) \\
& =n-\left(p t_{p}+q t_{q}\right)+(p+q)-\left(t_{p}+t_{q}\right)+1 \\
& =n+O\left(n^{3 / 4}\right) .
\end{aligned}
$$

This means that when $c d$ is known, the product $a b$ can be recovered with a Baby-step Giant-step type algorithm using $O\left(n^{3 / 8} / \sqrt{c d}\right)$ steps.

When $c=d$, we repeat the above analysis, and we find that we can recover $a b$ using $O\left(n^{3 / 8} / c\right)$ steps.

### 4.5.3 Decision Diffie-Hellman problem

The Decision Diffie-Hellman Problem was discussed in Section 2.3. We shall show that the Decision Diffie-Hellman problem is equivalent to a subgroup membership problem [10, 41], and the Computational Diffie-Hellman problem is equivalent to a splitting problem.

Let $G^{\prime}$ be a cyclic group with generator $g_{1}$. Let $G=G^{\prime} \times G^{\prime}$. Let $g_{2}$ be any element in $G^{\prime}$. Then the pair $\left(g_{1}, g_{2}\right)$ generates a non-trivial, proper subgroup $K$ of $G$. Let $H=\left\langle\left(1, g_{1}\right)\right\rangle$.

If $\Gamma^{\prime}$ is a family of groups, we can derive a family of group structures $\Gamma$ as above. For each group $G=G^{\prime} \times G^{\prime}$, there are $\left|G^{\prime}\right|$ different subgroups of the above form, and we sample from them with uniform probability. We get a subgroup membership problem $\mathcal{S} \mathcal{M}_{\Gamma}$ and splitting problem $\mathcal{S P}_{\Gamma}$.

Suppose we have a problem instance $(G, K, w)$. Now $K$ is defined by two elements $\left(g_{1}, g_{2}\right)$, and $w$ is defined by a pair $(x, y)$. This gives us a quadruple $\left(g_{1}, g_{2}, x, y\right)$. It is clear that $(x, y) \in K$ if an only if $\log _{g_{1}} g_{2} \log _{g} x \equiv \log _{g} y\left(\bmod \left|G^{\prime}\right|\right)$. The quadruple ( $g_{1}, g_{2}, x, y$ ) is a problem instance for the Decision Diffie-Hellman problem $\mathcal{D D H}_{G^{\prime}}$, and it is sampled with exactly the same probability. This means that $\mathcal{S M}_{\Gamma}$ is simply a reformulation of $\mathcal{D D H}_{\Gamma^{\prime}}$.

In the same way, if $(G, K, H, w)$ is an instance of $\mathcal{S} \mathcal{P}_{G}$, we get a quadruple $\left(g_{1}, g_{2}, x, y\right)$. The answer is two pairs $(x, z)$ and $\left(1, z^{\prime}\right)$, such that $\log _{g_{1}} g_{2} \log _{g_{1}} x \equiv$ $\log _{g_{1}} z\left(\bmod \left|G^{\prime}\right|\right)$. An instance of $\mathcal{C D} \mathcal{H}_{G^{\prime}}$ is a triple $\left(g_{1}, g_{2}, x\right)$, and the answer is $z$ as described.

It is clear that we can transform an instance of $\mathcal{S} \mathcal{P}_{(G, K, H)}$ into an instance of $\mathcal{C D H}_{G^{\prime}}$ by discarding $y$. We can transform an instance of $\mathcal{C D H}_{G^{\prime}}$ into an instance of $\mathcal{S P}{ }_{(G, K, H)}$ (where $K$ may vary) by sampling $y$ uniformly at random from $G^{\prime}$.

We see that $\mathcal{S P}_{\Gamma}$ is simply a reformulation of $\mathcal{C D} \mathcal{H}_{\Gamma^{\prime}}$.
It is worth noting that if $K$ is a subgroup of $G$ generated by $\left(g_{1}, g_{2}\right)$, and $K^{\prime}$ is generated by $\left(g_{1}^{a}, g_{1}^{b}\right)$, the map $(x, y) \mapsto\left(x^{a}, y^{b}\right)$ takes $K$ into $K^{\prime}$. The
map is an automorphism on $H$ when $b \not \equiv 0\left(\bmod \left|G^{\prime}\right|\right)$. This corresponds to the self-reducibility property of the Diffie-Hellman problems.

Sampling from $K=\left\langle\left(g_{1}, g_{2}\right)\right\rangle$ can be done by sampling $k$ uniformly at random from $\left\{0, \ldots,\left|G^{\prime}\right|-1\right\}$ and computing $\left(g_{1}, g_{2}\right)^{k}$. Sampling uniformly from $H$ can obviously be done by sampling uniformly from $G^{\prime}$.

Since $G$ is non-cyclic, Proposition 4.5 cannot be used to solve the splitting problem. But anyone who knows the discrete logarithm $\log _{g_{1}} g_{2}$ can compute the splitting, and this shows how to derive a trapdoor splitting problem.

### 4.5.4 Decision Composite Residuosity problem

This subgroup membership problem was first proposed by Paillier [35]. Our description also incorporates several later simplifications [14].

Let $n=p q$ be such that $\operatorname{gcd}(n, \phi(n))=1$. We shall consider the group $G=\mathbb{Z}_{n^{2}}^{*}$. It is clear that $|G|=\phi(n) n$.

Let $d \equiv n^{-1}(\bmod \phi(n))$. Then we can define a map from $\mathbb{Z}_{n}^{*}$ into $G$ by $[a]_{n} \mapsto\left[a^{d n}\right]_{n^{2}}$. It is clear that $a \equiv a^{d n}(\bmod n)$, so this map is one-to-one. It is also an homomorphism. We let the image of $\mathbb{Z}_{n}^{*}$ under this map be $K$.

Its inverse map $G \rightarrow \mathbb{Z}_{n}^{*}$ is the map $[a]_{n^{2}} \mapsto[a]_{n}$, often called reduction modulo $n$. It is clear that the elements $[1+a n]_{n^{2}}$ map to $[1]_{n}$, and that any element that maps to $[1]_{n}$ is of this form. We let $H$ be this kernel of reduction modulo $n$.

It is quite clear that $K \cap H=\{1\}$ and that $K H=G$.
Now we note that

$$
(1+n)^{a} \equiv 1+\binom{a}{1} n+\binom{a}{2} n^{2} \cdots \equiv 1+a n \quad\left(\bmod n^{2}\right)
$$

From this, we see that discrete logarithms are easy to compute in $H$. One group isomorphism $\mathbb{Z}_{n} \rightarrow H$ is the map $[a]_{n} \mapsto[1+n]_{n^{2}}^{a}$. Since for any $[b]_{n^{2}} \in H, b \equiv 1$ $(\bmod n)$, the reverse isomorphism is $[b]_{n^{2}} \mapsto[(b-1) / n]_{n}$.

The subgroup membership problem $\mathcal{S M}_{(G, K)}$ is called the Decision Composite Residuosity problem $\left(\mathcal{D C R}_{n}\right)$. The corresponding splitting problem $\mathcal{S P}_{(G, K, H)}$ is called the Computational Composite Residuosity problem $\left(\mathcal{C C R}_{n}\right)$. This is equivalent to the subgroup discrete logarithm problem $\mathcal{S D}_{(G, K, H,[1+n])}$.

We can sample elements uniformly from $K$ and $H$ by sampling elements uniformly in $\mathbb{Z}_{n}^{*}$ and $\mathbb{Z}_{n}$, and then using the corresponding homomorphisms.

It is quite clear that anyone who can solve $\mathcal{F A C} \mathcal{T}_{n}$ can solve $\mathcal{C C R}_{n}$. Likewise, anyone who can solve $\mathcal{R S} \mathcal{A}_{n}$ with encryption exponent $n$ can solve $\mathcal{C C} \mathcal{R}_{n}$. It is not known if $\mathcal{C C R}$ is equivalent to either of these, but the best known method for solving the $\mathcal{C C R}{ }_{n}$ and $\mathcal{D C} \mathcal{R}_{n}$ is to factor the modulus $n$.

Damgård and Jurik [14] describe a generalised group structure, using $\mathbb{Z}_{n^{s+1}}^{*}$ for some $s \geq 1$. As above, they show that $G=\mathbb{Z}_{n^{s+1}}^{*} \simeq \mathbb{Z}_{n}^{*} \times \mathbb{Z}_{n^{s}} \simeq K \times H$. They get a subgroup membership problem $\mathcal{S M}_{(G, K)}$, and they show that it does not get much easier as $s$ increases. We sketch the proof.

Let $K_{s}$ and $H_{s}$ be the subgroups of $\mathbb{Z}_{n^{s+1}}^{*}$ isomorphic to $\mathbb{Z}_{n}^{*}$ and $\mathbb{Z}_{n^{s}}$, respectively. Note that $[1+n]_{n^{s+1}}$ is a generator for $H_{s}$. We shall proceed by induction, using the following diagram:


Consider the map $f: \mathbb{Z}_{n^{2}}^{*} \rightarrow \mathbb{Z}_{n^{s+1}}$ given by $f\left([w]_{n^{2}}\right)=\left[w^{n^{s-1}}\right]_{n^{s+1}}$. Since for any $a, b \in \mathbb{Z},\left(a+b n^{t}\right)^{n^{s+1-t}} \equiv a^{n^{s+1-t}}\left(\bmod n^{s+1-t}\right)$, and composition with reduction modulo $n$ shows that $f\left(\mathbb{Z}_{n^{2}}\right) \subseteq \mathbb{Z}_{n^{s+1}}$, it is clear that $f$ is a well-defined group homomorphism.

So suppose $x \in H_{1}$. Then for some $a \in\{0, \ldots, n-1\}, x=\left[(1+n)^{a}\right]_{n^{2}}$ and $f(x)=\left[\left((1+n)^{a}\right)^{n^{s-1}}\right]_{n^{s+1}}=\left[(1+n)^{a n^{s-1}}\right]_{n^{s+1}} \in H_{s}^{n^{s-1}}$. This means that $f\left(H_{1}\right)=H_{s}^{n^{s-1}}$.

For any element $x \in K_{1}$, there must be an integer $w$ such that $x=\left[w^{n}\right]_{n^{2}}$. But then $f(x)=\left[\left(w^{n}\right)^{n^{s-1}}\right]_{n^{s+1}}=\left[w^{n^{s}}\right]_{n^{s+1}}$, so $f\left(K_{1}\right) \subseteq K_{s}$. It is clear that $f$ is a group homomorphism. Composing with reduction modulo $n$, and noting that $x \mapsto x^{n}$ is an automorphism of $\mathbb{Z}_{n}^{*}$, we see that the kernel of $f$ restricted to $K_{1}$ is trivial, so $f\left(K_{1}\right)=K_{s}$.

By Proposition 4.3, this means that if we can distinguish $\mathcal{S M}_{\left(K_{s} H_{s}^{n s-1}, K_{s}\right)}$, then we can distinguish $\mathcal{S M}_{\left(K_{1} H_{1}, K_{a}\right)}$.

Next, we consider the homomorphism $f^{\prime}: K_{s} H_{s} \rightarrow K_{s-1} H_{s-1}$ given by $f^{\prime}\left([w]_{n^{s+1}}\right)=[w]_{n^{s}}$. It is clear that the kernel of $f^{\prime}$ is $H_{s}^{n^{s-1}}$. Also, if $w \equiv w^{\prime}$ $\left(\bmod n^{s}\right)\left(\operatorname{and} \operatorname{gcd}(w, n)=\operatorname{gcd}\left(w^{\prime}, n\right)=1\right)$, then $w \equiv w^{\prime}(1+n)^{r n^{s-1}}\left(\bmod n^{s+1}\right)$ for some $r \in \mathbb{Z}$.

Let $A$ be an algorithm that takes as input $[w]_{n^{s}} \in K_{s-1} H_{s-1}$, samples $r$ uniformly from $\{0, \ldots, n-1\}$, and outputs $\left[w(1+n)^{r n^{s-1}}\right]_{n^{s+1}}$. It is quite clear that if the input is uniformly distributed in $K_{s-1}$, then the output is uniformly distributed in $K_{s} H_{s}^{n^{s-1}}$. Likewise, if $x \in H_{s-1}$ is uniformly distributed in $H_{s-1}$, then the output is almost uniformly distributed in $H_{s}$.

Again, by Proposition 4.3, if we can distinguish $\mathcal{S M}_{\left(K_{s} H_{s}, K_{s} H_{s}^{n s-1}\right)}$, then we can distinguish $\mathcal{S} \mathcal{M}_{\left(K_{s-1} H_{s-1}, K_{s-1}\right)}$.

Now we consider only distinguishing algorithms up to a certain fixed cost. Let $\epsilon_{i}$ be the maximal advantage of these algorithms against the subgroup membership problem $\mathcal{S} \mathcal{M}_{\left(K_{i} H_{i}, K_{i}\right)}$. By Proposition 4.2 and the above discussion, we
get that $\epsilon_{i} \leq \epsilon_{i-1}+3 \epsilon_{1} / 2$, or $\epsilon_{s} \leq(3 s-1) \epsilon_{1} / 2$. In other words, the advantage increases at most linearly as $s$ increases.

Galbraith [20] gave an elliptic curve generalisation of this group structure. The basic idea is that modulo $n^{2}$, the projective point $(n: 1: 0)$ has order $n$, but modulo $n$, it reduces to the point at infinity. So just as for $\mathbb{Z}_{n^{2}}^{*}, E\left(\mathbb{Z}_{n^{2}}\right) \simeq$ $E\left(\mathbb{Z}_{n}\right) \times \mathbb{Z}_{n}$, with $\mathbb{Z}_{n}$ taken additively.

It is easy to sample elements uniformly at random from the part isomorphic to $\mathbb{Z}_{n}$, but to sample from the part isomorphic to $E\left(\mathbb{Z}_{n}\right)$, we need a generator for $E\left(\mathbb{Z}_{n}\right)$ (or a subgroup, if $E\left(\mathbb{Z}_{n}\right)$ is non-cyclic).

### 4.5.5 Okamoto-Uchiyama

A precursor to the Composite Residuosity group structure was described by Okamoto and Uchiyama [34].

Let $p$ and $q$ be primes such that $\operatorname{gcd}(p q, \phi(p q))=1$, and let $n=p^{2} q$. As above, we find that $\mathbb{Z}_{n}^{*} \simeq \mathbb{Z}_{p}^{*} \times \mathbb{Z}_{q}^{*} \times \mathbb{Z}_{p}$, where $\mathbb{Z}_{p}$ is taken to be an additive group. $K$ is the subgroup isomorphic to $\mathbb{Z}_{p}^{*} \times \mathbb{Z}_{q}^{*}$, and $H$ is the subgroup isomorphic to $\mathbb{Z}_{p}$.

We can sample elements uniformly from $K$ by sampling uniformly from $\mathbb{Z}_{n}^{*}$, and then computing the $n$th power. Since $\operatorname{gcd}(p q, \phi(p q))=1$, raising to the $n$th power is an automorphism on $K$, while it is the zero homomorphism on $H$.

Sampling elements from $H$ is a bit more tricky. Any integer representing a residue class in $H$ is congruent to 1 modulo $p$. This means that anyone who knows a non-trivial element $[a]$ in $H$ can factor $n$ by computing $\operatorname{gcd}(a-1, n)$.

But given an element $g$ whose order is divisible by $a b p$, we can sample elements almost uniformly at random from $G \backslash K$ by sampling $u$ from $\{0, \ldots, n-1\}$, subject to $\operatorname{gcd}(u, n)=1$. Then we sample $r \in \mathbb{Z}_{n}^{*}$ uniformly at random and compute $r^{n} g^{u}$.

Anyone who can solve the subgroup discrete logarithm problem, can recover the order of $H$, which is $p$, so he can factor $n$. Anyone who can factor $n$ can solve the splitting problem. As discussed, anyone who can solve the subgroup splitting problem can factor $n$ and then solve the discrete logarithm problem in $H$. So the subgroup discrete logarithm problem and the splitting problem are equivalent, and both are equivalent to factoring the underlying modulus.

### 4.5.6 Groups of known order

Let $p$ and $q$ be primes, and let $n=p q$. Let $G$ be an abelian group of order $n$. Then $G$ has unique subgroups $K, H$ of order $p$ and $q$, respectively. If the order of $G$ is easily computable, but the orders of $K$ and $H$ are hard to find, then we have a potential candidate for a subgroup membership problem.

We shall describe two families of such groups. Note that for both families, the two subgroups $K$ and $H$ are essentially the same, so we get symmetric subgroup membership problems.

Let $n$ be a composite number with two prime factors such that $2 n+1$ is also prime. Then the multiplicative group $\mathbb{F}_{2 n+1}^{*}$ has order $2 n$, and we have a unique subgroup of order $n$. This group structure was suggested in [31]. A related topic was discussed in [26]. We denote the derived symmetric subgroup membership problem by $\mathcal{S O U}_{n}$.

Let $n$ be a composite number with two prime factors, let $p^{\prime}$ be a prime such that $\left|p^{\prime}+1-n\right|<2 \sqrt{p^{\prime}}$, and let $E$ be an elliptic curve defined over $\mathbb{F}_{p^{\prime}}$ such that $\# E\left(\mathbb{F}_{p^{\prime}}\right)=n$. We derive a symmetric subgroup membership problem $\mathcal{S O} \mathcal{U}_{E / \mathbb{F}_{p^{\prime}}}$. (One would probably resort to complex multiplication techniques to find $E$.)

To sample elements almost uniformly from the subgroups $K$ and $H$, generators for $K$ and $H$ must be made available.

It seems plausible that $\mathcal{S O U}_{n}$ and $\mathcal{S O U}_{E / \mathbb{F}_{p^{\prime}}}$ both are hard. Deciding if an element is in one of the subgroups is equivalent to deciding if its order is small or large. It is clear that anyone who can factor the modulus $n$ can also break the system. Likewise, anyone who can compute discrete logarithms in the subgroups $K$ and $H$ can factor the modulus and hence break the system. Anyone who can compute discrete logarithms in $G$, but not in $K$ and $H$, can decide the subgroup membership problems. Also, anyone who can decide if an element is a generator for the whole group, can decide the subgroup membership problems.

It is possible for $n$ to have more than two prime factors, or for the two prime factors to be of different size. While factoring algorithms like the number field sieve does not seem able to take advantage of such a structure, the elliptic curve method [25] does, and this gives a fairly large lower bound on the size of the smallest prime factor.

### 4.5.7 Groups of unknown order

Let $p=2 a c+1$ and $q=2 b d+1$ be $\tau$-bit primes such that $a, b, c$ and $d$ are pairwise relatively prime, and let $n=p q$. Let $G=Q_{n}$, let $K$ be the unique subgroup of order $2 a b$, and let $H$ be the unique subgroup of order $c d$. Then we get a splitting problem $\mathcal{S P}_{(G, K, H)}$ and a symmetric subgroup membership problem $\mathcal{S S M}_{(G, K, H)}$ which we denote by $\mathcal{G O \mathcal { U } _ { n }}$.

In order to sample elements almost uniformly from $K$ and $H$, generators for both groups must be made public and Lemma 2.4 applied.

In general, the factors $a, b, c$ and $d$ should be large. Almost certainly, $|K|$ is significantly smaller than $n$, so if the product $a b=|K|$ is known, it will be feasible to find $a$ and $b$ via factorisation, and then factor $n$. Therefore, it should be difficult to compute $|K|$, even when a generator for $K$ is known. The same holds for $H$.

This means that $a, b, c$ and $d$ should each contain at least one big prime factor. This prevents anyone from finding the order of $K$ and $H$ by computing discrete
logarithms. Also, it blocks Pollard $\rho$-type attacks on the subgroup, which would have been feasible if $|K|$ or $|H|$ were small.

Note that unlike in the previous section, the order of $G$ is not known, so the prime factors in $a, b, c$ and $d$ can still be quite small.

As in Section 4.5.2, we can let $c=d$. Then $K$ is the subgroup of order $a b$ and $H$ is some subgroup of order $c$. Since $c \mid(n-1), c$ becomes known, but it does not seem to matter, even if an element of order $c$ is made public. $c$ should not be too large, however, as the analysis in Section 4.5.2 also applies to this case. And since $c$ is known, the subgroup membership problem $\mathcal{S M}_{(G, H)}$ is trivial.

We denote this subgroup membership problem by $\mathcal{G O} \mathcal{U}_{n}^{\prime}$.

### 4.6 Further reductions

Let $p=2 a c+1, q=2 b d+1, p \equiv q \equiv 3(\bmod 4)$, and $n=p q$, just as in Sections 4.5.1 and 4.5.7. Let $J_{n}$ be the subgroup $\mathbb{Z}_{n}^{*}$ of elements with Jacobi symbol $1, Q_{n}$ be the subgroup of quadratic residues, $K$ be the subgroup of order $a b$ and $H$ the subgroup of order $c d$. We denote the subgroup of order 2 generated by -1 by $T$. We get the following subgroup lattice:


We have already discussed the subgroup membership problems $\mathcal{S M}_{\left(J_{n}, Q_{n}\right)}$, $\mathcal{S} \mathcal{M}_{\left(Q_{n}, K\right)}$ and $\mathcal{S} \mathcal{M}_{\left(Q_{n}, H\right)}$. We first discuss some combinations of these problems.

Proposition 4.13. Suppose there is a distinguisher for $\mathcal{S M}_{\left(J_{n}, K\right)}$ with advantage $\epsilon$. Then either there is a distinguisher for $\mathcal{S} \mathcal{M}_{\left(J_{n}, Q_{n}\right)}$ with advantage at least $\epsilon / 3$, or there is a distinguisher for $\mathcal{S M}_{\left(Q_{n}, K\right)}$ with advantage at least $\epsilon / 3$.
Proof. Apply Proposition 4.2.
Proposition 4.14. For any distinguisher for $\mathcal{S M}_{\left(J_{n}, T K\right)}$ with advantage $\epsilon$, there is a corresponding distinguisher for $\mathcal{S}_{\left(Q_{n}, K\right)}$ with the same advantage, and vice versa.

Proof. We shall show how to apply Proposition 4.3 to both directions.
First, we note that the map from $J_{n}$ to $Q_{n}$ given by $x \mapsto x^{2}$ is an automorphism on the groups $Q_{n}, K$ and $H$. If $X$ is uniformly distributed on $J_{n} \backslash T K$, then the projection of $X$ on $H$ is uniformly distributed on $H \backslash\{1\}$. Since squaring
is an automorphism, the projection of $X^{2}$ on $H$ is also uniformly distributed on $H \backslash\{1\}$.

If $X$ is uniformly distributed on $T K$, then $X^{2}$ is clearly uniformly distributed on $K$. So by Proposition 4.3 we have shown that $\mathcal{S M}_{\left(J_{n}, T K\right)}$ is not more difficult than $\mathcal{S M}_{\left(Q_{n}, K\right)}$.

Second, let $Y$ be uniformly distributed on $\{0,1\}$. If $X$ is uniformly distributed on $Q_{n} \backslash K$, then $X(-1)^{Y}$ is uniformly distributed on $J_{n} \backslash T K$. If $X$ is uniformly distributed on $K$, then $X(-1)^{Y}$ is uniformly distributed on $T K$. So again by Proposition 4.3 we have shown that $\mathcal{S M}_{\left(Q_{n}, K\right)}$ is not more difficult than $\mathcal{S M}_{\left(J_{n}, T K\right)}$.

We shall now combine the above group structure with the group structure of Section 4.5.4. Let $J_{n^{2}}$ be the elements of $\mathbb{Z}_{n^{2}}^{*}$ with Jacobi symbol 1, and let $Q_{n^{2}}$ be the quadratic residues of $\mathbb{Z}_{n^{2}}^{*}$.

Recall that we have a homomorphism $\mathbb{Z}_{n^{2}}^{*} \rightarrow \mathbb{Z}_{n}^{*}$ given by taking the residue class of $\mathbb{Z}_{n^{2}}$ represented by the integer $w$ to the residue class of $\mathbb{Z}_{n}$ represented by $w$. Let $H^{\prime}$ be the kernel of this homomorphism.

We also have a homomorphism $\mathbb{Z}_{n}^{*} \rightarrow \mathbb{Z}_{n^{2}}^{*}$ given by taking the residue class of $\mathbb{Z}_{n}^{*}$ represented by the integer $w$ to the residue class of $\mathbb{Z}_{n^{2}}^{*}$ represented by $w^{n}$. Let $K^{\prime}$ be the image of $\mathbb{Z}_{n}^{*}$ under this homomorphism. We identify the subgroups of $\mathbb{Z}_{n}^{*}$ with the corresponding subgroups of $K^{\prime}$.

We get the following subgroup lattice:


Proposition 4.15. For any distinguisher for $\mathcal{S M}_{\left(Q_{n^{2}}, Q_{n}\right)}$ with advantage $\epsilon$, there is a corresponding distinguisher $\mathcal{S} \mathcal{M}_{\left(\mathbb{Z}_{n 2}^{*}, K^{\prime}\right)}$ with the same advantage.

Suppose an element $j \in \mathbb{Z}_{n^{2}}^{*} \backslash J_{n^{2}}$ is known. Then for any distinguisher for $\mathcal{S} \mathcal{M}_{\left(\mathbb{Z}_{n}^{*}, K^{\prime}\right)}$ with advantage $\epsilon$, there is a corresponding distinguisher $\mathcal{S M}_{\left(Q_{n}, Q_{n}\right)}$ with the same advantage.
Proof. Obviously, the map $x \mapsto x^{2}$ takes the uniform distributions on $\mathbb{Z}_{n^{2}}^{*} \backslash K^{\prime}$ and $K^{\prime}$ to the uniform distributions on $Q_{n^{2}} \backslash Q_{n}$ and $Q_{n}$, respectively. This means
that $\mathcal{S M}_{\left(\mathbb{Z}_{n^{2}}^{*}, K^{\prime}\right)}$ cannot be more difficult than $\mathcal{S} \mathcal{M}_{\left(Q_{n^{2}}, Q_{n}\right)}$, by Proposition 4.3, and the first claim follows.

Let $Y$ and $Z$ be uniformly distributed on $\{0,1\}$. If $X$ is uniformly distributed on $Q_{n}$, then $X(-1)^{Y} j^{n Z}$ is uniformly distributed on $K^{\prime}$. If $X$ is uniformly distributed on $Q_{n^{2}} \backslash Q_{n}$, then $X(-1)^{Y} j^{n Z}$ is uniformly distributed on $\mathbb{Z}_{n^{2}}^{*} \backslash K^{\prime}$. As above, we find that $\mathcal{S} \mathcal{M}_{\left(Q_{n^{2}}, Q_{n}\right)}$ cannot be more difficult than $\mathcal{S} \mathcal{M}_{\left(\mathbb{Z}_{n 2}^{*}, K^{\prime}\right)}$, by Proposition 4.3, and the second claim follows.

Proposition 4.16. Suppose there is a distinguisher for $\mathcal{S M}_{\left(Q_{n^{2}}, K\right)}$ with advantage $\epsilon$. Then either there is a distinguisher for $\mathcal{S M}_{\left(Q_{n^{2}}, Q_{n}\right)}$ with advantage $\epsilon / 3$, or there is a distinguisher for $\mathcal{S} \mathcal{M}_{\left(Q_{n}, K\right)}$ with advantage $\epsilon / 3$.

Proof. Apply Proposition 4.2.

## Chapter 5

## Homomorphic cryptosystems

In this section, we shall study three different homomorphic cryptosystems. We shall prove all three systems semantically secure against chosen plaintext attacks based on subgroup membership problems. All three cryptosystems are homomorphic with respect to a group operation. For several of the cryptosystems, the group operation can be regarded as addition modulo an integer, which may be very convenient.

We shall also discuss the security of some of the cryptosystems against a non-adaptive chosen ciphertext attack, under a non-standard assumption.

### 5.1 The general cryptosystem

Let $\mathcal{T S P}_{\Gamma}$ be a trapdoor splitting problem, and suppose that to every tuple ( $G, K, H, \sigma$ ), there is an algorithm for sampling $K$ (almost) uniformly at random, an associated group $G^{\prime}$, and a probability space $F$ on a set of easy to compute maps $\left\{f: G^{\prime} \rightarrow G\right\}$. We require that for every map $f$ in $F$, the composition $\pi_{H} \circ f: G^{\prime} \rightarrow H$ is a group isomorphism, and it is easy to derive a group isomorphism $f^{\prime}: H \rightarrow G^{\prime}$ such that $f^{\prime}=\left(\pi_{H} \circ f\right)^{-1}$.

The basic idea is that $H$ carries the message, and $K$ contains "noise" that hides the message. $f$ is used to insert the message in $H$. The splitting algorithm can be used to remove the noise, and $f^{\prime}$ can be used to recover the message. The exact construction of $\Pi_{1}$ from $\mathcal{T S} \mathcal{P}_{\Gamma}$ is described in Figure 5.1.

Proposition 5.1. The cryptosystem $\Pi_{1}$ is an homomorphic public key system.
Proof. It is quite obvious that $\mathcal{K}, \mathcal{E}$ and $\mathcal{D}$ are efficient, and that $\mathcal{D}$ is deterministic. It is also clear that if $r f(m)$ is the ciphertext output by $\mathcal{E}$, then the projection on $H$ is $\pi_{H}(f(m))$, so $\mathcal{D}$ will output $m$. Therefore, $\Pi_{1}$ is a public key system.

## Key generation.

Input: $\mathcal{T S}_{\Gamma}, \tau$.

1. $(G, K, H, \sigma, F) \leftarrow \Gamma_{\tau}$.
2. $f \leftarrow F$.
3. $p k \leftarrow\left(G, K, H, G^{\prime}, f\right)$.
4. $s k \leftarrow\left(G, \sigma, G^{\prime}, f^{\prime}\right)$.
5. Output ( $p k, s k$ ).

Output: A public-private key pair $(p k, s k)$.

## Encryption.

Input: $p k=\left(G, K, H, G^{\prime}, f\right), m \in G^{\prime}$.

1. $x \leftarrow f(m), r \approx K$.
2. $c \leftarrow r x$.
3. Output $c$.

Output: A ciphertext $c \in G$.

## Decryption.

Input: $s k=\left(G, \sigma, G^{\prime}, f^{\prime}\right), c \in G$.

1. $\left(c_{1}, c_{2}\right) \leftarrow \sigma(c)$.
2. $m \leftarrow f^{\prime}\left(c_{2}\right)$.
3. Output $m$.

Output: A message $m \in G^{\prime}$.

Figure 5.1: The cryptosystem $\Pi_{1}$.

Finally, if $r f(m)$ and $r^{\prime} f\left(m^{\prime}\right)$ are two ciphertexts, then the decryption of $r f(m) r^{\prime} f\left(m^{\prime}\right)$ is clearly $f^{\prime}\left(\pi_{H}\left(r r^{\prime} f(m) f\left(m^{\prime}\right)\right)\right)=f^{\prime}\left(\pi_{H}(f(m))\right) f^{\prime}\left(\pi_{H}\left(f\left(m^{\prime}\right)\right)\right)=$ $\mathrm{mm}^{\prime}$, and the cryptosystem is homomorphic.

Proposition 5.2. Suppose that for any map $f \in F$, we have that for all $m \in G^{\prime}$, $f(m)=\pi_{H}(f(m))$. Then the public key cryptosystem $\Pi_{1}$ is one-way if and only if the splitting problem $\mathcal{S P}_{\Gamma}$ is hard.

Proof. We prove the statements by giving reductions.
It is clear that any algorithm that solves the splitting problem can be used instead of the trapdoor function $\sigma$ used in $\mathcal{D}$.

Now suppose that $(G, K, H, x)$ is an instance of a splitting problem. The group $G^{\prime}$ is given, and we can sample $f$ from $F$, to construct a public key $\left(G, K, H, G^{\prime}, f\right)$, which has been sampled just as $\mathcal{K}$ would have done it.

Note that when messages are sampled uniformly from $G^{\prime}, \mathcal{E}$ samples ciphertexts uniformly from $G$. Therefore, $x$ is a ciphertext for $p k$ sampled exactly as normal ciphertexts, and if $m$ is the decryption of $x$, then $f(m)$ is the projection on $H$ of $x$, and $\left(x(f(m))^{-1}, f(m)\right)$ is the splitting of $x$.

Let $\mathcal{S P}_{(G, K, H)}$ be a splitting problem such that $H$ is cyclic. Suppose there is some efficient, deterministic algorithm to compute discrete logarithms in $H$. There is a natural map $\mathbb{Z} \rightarrow G$ given by $m \mapsto g^{m}$, where $g \in G$ is such that its projection on $H$ is a generator for $H$. If we decide on some fixed map $\kappa: \mathbb{Z}_{|H|} \rightarrow \mathbb{Z}$ to choose representatives for the residue classes of $\mathbb{Z}_{|H|}$, we get a map $f: \mathbb{Z}_{|H|} \rightarrow G$ given by $[m] \mapsto g^{\kappa(m)}$.

It is quite clear that if we have a trapdoor subgroup discrete logarithm problem $\mathcal{T S D} \mathcal{L}_{\Gamma}$, then we can use this map to get a variant cryptosystem $\Pi_{1^{\prime}}$. It is described in Figure 5.2.

For $\Pi_{1^{\prime}}$, Proposition 5.2 is not true unless the element $g$ is a generator $H$. We can recover the following, slightly weaker result.

Proposition 5.3. The public key cryptosystem $\Pi_{1^{\prime}}$ is one-way if and only if the subgroup discrete logarithm problem $\mathcal{S D}_{\mathcal{L}_{\Gamma}}$ is hard.

Proof. We prove the statements by giving reductions.
It is clear that any algorithm that solves the subgroup discrete logarithm problem can be used instead of the trapdoor function $\sigma$ and the logarithm computation used in $\mathcal{D}$.

Let $(G, K, H, g, y)$ be a problem instance sampled from a subgroup discrete logarithm problem $\mathcal{S D} \mathcal{L}_{\Gamma}$. We get the public key $p k=(G, K, H, g)$, and this has been sampled exactly according to $\mathcal{K}(\tau) . y$ is a ciphertext under the public key $p k$, and it has been sampled exactly as ciphertexts are sampled. Its decryption $m$ is congruent to $\log _{(H, g)} y$ modulo $|H|$, so anyone who can decrypt the ciphertext can solve the subgroup discrete logarithm problem.

## Key generation.

Input: $\mathcal{T S D}_{\Gamma}, \tau$.

1. $(G, K, H, \sigma, g) \leftarrow \Gamma_{\tau}$.
2. $p k \leftarrow(G, K, H, g)$.
3. $\left(g_{1}, g_{2}\right) \leftarrow \sigma(g)$.
4. $s k \leftarrow\left(G, \sigma, g_{2}\right)$.
5. Output ( $p k, s k$ ).

Output: A public-private key pair ( $p k, s k$ ).

## Encryption.

Input: $p k=(G, K, H, g),[m] \in \mathbb{Z}_{|H|}$.

1. $x \leftarrow g^{m}, r \approx K$.
2. $c \leftarrow r x$.
3. Output $c$.

Output: A ciphertext $c \in G$.

## Decryption.

Input: $s k=\left(G, \sigma, G^{\prime}, f^{\prime}\right), c \in G$.

1. $\left(c_{1}, c_{2}\right) \leftarrow \sigma(c)$.
2. $m \leftarrow \log _{g_{2}} c_{2}$.
3. Output $m$.

Output: A message $m \in \mathbb{Z}_{|H|}$.

Figure 5.2: The cryptosystem $\Pi_{1^{\prime}}$.

Now we use the indistinguishability property of the subgroup membership problem to prove a stronger security result for the cryptosystems $\Pi_{1}$ and $\Pi_{1^{\prime}}$.

Theorem 5.4. The public key cryptosystems $\Pi_{1}$ and $\Pi_{1^{\prime}}$ are semantically secure if and only if the subgroup membership problem $\mathcal{S M}_{\Gamma}$ is hard.
Proof. Again, we give reductions.
Let $\left(G, K, H, G^{\prime}, f\right)$ be a public key. Suppose we have a ciphertext and want to decide if $c$ decrypts to a message $m$ or to some other, random message. First, we set $x=f(m)$. If $\mathcal{D}(s k, c)=m$, then $c x^{-1}$ is uniformly distributed in $K$. If $\mathcal{D}(s k, c) \neq m$, then $c x^{-1}$ is uniformly distributed in $G \backslash K$. Hence, we can decide if $c$ decrypts to $m$ or not with the same advantage that we can decide the subgroup membership problem $\mathcal{S M}_{(G, K)}$.

Next, suppose we have a subgroup membership problem $\mathcal{S M}_{(G, K)}$ and an instance $x$. As in the proof of Proposition 5.2, we first sample a public key. Let $m \in p k_{M}$ be a message, and sample $x^{\prime}$ uniformly from $K$. If $x \in K$, then $x x^{\prime} f(m)$ is an encryption of $m$. If $x \in G \backslash K$, then $x x^{\prime} f(m)$ is not an encryption of $m$, but of some other, random message. In both cases, the distribution of the ciphertexts are exactly as $\mathcal{E}$ would produce, so we can decide the subgroup membership problem $\mathcal{S} \mathcal{M}_{(G, K)}$ with the same advantage that we can distinguish encryptions of $m$ from other encryptions of other messages.

We briefly describe various concrete instances of $\Pi_{1}$ and $\Pi_{1^{\prime}}$.
Goldwasser-Micali We can instantiate $\Pi_{1}$ with the Quadratic Residue problem from Section 4.5.1 (using a suitable family of moduli), and we get the Goldwasser-Micali cryptosystem [22].

In this case, the public key is $p k=(n)$. Sampling an element uniformly at random is done by first sampling an element uniformly at random from $\mathbb{Z}_{n}^{*}$, then squaring it. The message space is $\{1,-1\}$, which is a subgroup of $J_{n}$, and $f=f^{\prime}$ is the identity.

The secret key is one of the prime factors of $n$, say $s k=(p)$. To decrypt a ciphertext $c$, one simply computes the Legendre symbol of $c$ with respect to $p$.

If we let the message space be $\{0,1\}$ with the mapping $0 \mapsto 1$ and $1 \mapsto-1$, then the cryptosystem becomes additive and we get an instance of $\Pi_{1^{\prime}}$.

ElGamal We can instantiate $\Pi_{1}$ with the Computational Diffie-Hellman problem from Section 4.5.3 (using a suitable family of groups), we get the classical ElGamal cryptosystem [18].

In this case, the public key is $p k=\left(G^{\prime}, g, y\right)$ where $G^{\prime}$ is the underlying group structure, $g$ is a generator, and $(g, y)$ defines the subgroup $K$. The subgroup $H$ is then trivially isomorphic to $G^{\prime}$, which is the message space. Sampling uniformly in $K$ is simply a matter of sampling $k$ uniformly from $\left\{0, \ldots,\left|G^{\prime}\right|-1\right\}$ and computing $(g, y)^{k}=\left(g^{k}, y^{k}\right)$.

The secret key is $s k=\left(G^{\prime}, a\right)$, where $a$ is an integer such that $y=g^{a}$. Then the splitting of $\left(z, z^{\prime}\right)$ can be computed as $\left(\left(z, z^{a}\right),\left(1, z^{\prime} z^{-a}\right)\right)$.

Paillier The Computational Composite Residuosity problem in Section 4.5.4, $\mathcal{C C R}_{n}$, can be considered either as a splitting problem or as a subgroup discrete logarithm problem. We take it to be the latter, so the Paillier cryptosystem [35] is an instance of $\Pi_{1^{\prime}}$.

The public key will be $p k=(n)$, and the secret key $s k=(d)$, where $d$ is the order of $K$, and the splitting is computed using Proposition 4.5.

Paillier [35] also suggested a subgroup variant. This is essentially the same system, but the modulus is chosen as in Section 4.5.7 and $K$ is taken to be a suitable subgroup of the $n$th powers (which are isomorphic to $\mathbb{Z}_{n}^{*}$ ). Propositions 4.15 and 4.16 say that the subgroup variant is secure if both $\mathcal{D C} \mathcal{R}_{n}$ and $\mathcal{G O} \mathcal{U}_{n}$ are hard.

Okamoto-Uchiyama The construction in Section 4.5.5 is a subgroup discrete logarithm problem, and we get an instance of $\Pi_{1^{\prime}}$ [34].

Let $p=2 a+1, q=2 b+1$ be primes and $n=p^{2} q$. The public key consists of $(n, g)$, where $g \in \mathbb{Z}_{n}^{*}$ has order divisible by $p$.

An interesting point is that given only the modulus $n$, we can sample a public key with almost the same distribution as the key generation algorithm. To see this, note that a random element in $\mathbb{Z}_{n}^{*}$ has order divisible by $p$ with probability approximately $1-1 / p$.

Now suppose that we know that $p$ is a $t$-bit integer, that is, $2^{t-1} \leq p<2^{t}$. If we sample $a$ from $\left\{2^{t}, \ldots, 2^{t+1}-1\right\}$, then $p<a<4 p$. We sample $r$ uniformly at random from $\mathbb{Z}_{n}^{*}$ and create the ciphertext $r^{n} g^{a}$. This is a valid ciphertext with decryption $m$. So we know that $a \equiv m(\bmod p)$, or that $a-m=k p$. Since $0 \leq m<p, k$ must be 1,2 or 3 . In other words, if we can decrypt ciphertexts for this instance of the cryptosystem, we can factor the underlying modulus.

Naccache-Stern If we instantiate $\Pi_{1^{\prime}}$ with $\mathcal{H} \mathcal{R}_{n}$, we get the Naccache-Stern cryptosystem [29]. By instantiating $\Pi_{1^{\prime}}$ with $\mathcal{H} \mathcal{R}_{n}^{\prime}$, we get a similar variant, which fits into both $\Pi_{1}$ and $\Pi_{1^{\prime}}$, but the latter is most natural.

Let $p=2 a c+1, q=2 b d+1$ be primes, such that $a, b$ are primes and $c, d$ are smooth integers (either equal, or relatively prime). Let $n=p q$. The public key is $(n, g)$, where $g \in G$ defines the function $f$ by $f(m)=g^{m}$. If $\operatorname{gcd}(c, d)=1$, then $g$ should be an element of order divisible by $a b c d$. If $c=d$, then $g$ should be an element of order $c$ that is not congruent to 1 modulo either $p$ or $q$.

Combinations It is possible to combine $\mathcal{H} \mathcal{R}_{n}$ (or $\mathcal{H R}_{n}^{\prime}$ ) with $\mathcal{C C R}_{n}$ to get an instance of $\Pi_{1^{\prime}}$. Proposition 4.15 and Proposition 4.16 say that the resulting instance of $\Pi_{1^{\prime}}$ is secure if both $\mathcal{H} \mathcal{R}_{n}$ (or $\mathcal{H} \mathcal{R}_{n}^{\prime}$ ) and $\mathcal{C C} \mathcal{R}_{n}$ are hard. Such a cryptosystem would have somewhat higher bandwidth. The disadvantages would include a larger public key (one would need to store generators for $K$ and $G$ ) and slower decryption.

### 5.2 Non-standard assumptions

To motivate the work in this section, we briefly describe a non-adaptive chosen ciphertext attack against the Paillier cryptosystem which may be faster than factoring the modulus. This is basically a variant of the attacks in [15], but using a trick of looking for smooth numbers in $\mathbb{Z}_{n}$, not in $\mathbb{Z}_{n^{2}}$.

So let $B$ be a set of small primes (possibly also including -1 ). Note that every number in $B$ gives us a valid ciphertext for the Paillier cryptosystem. We use the decryption oracle to obtain a decryption of every number in $B$, and store them in a table.

When we get the challenge ciphertext $c$, we choose random $n$th powers $c^{\prime}$, which decrypt to 0 , until $c c^{\prime} \bmod n$ is $B$-smooth. If the factorisation of $c c^{\prime}$ is $\prod_{\ell_{i} \in B} \ell_{i}^{\alpha_{i}}$, we can create a ciphertext $c^{\prime \prime}=c c^{\prime} \prod_{\ell_{i} \in B} \ell_{i}^{-\alpha_{i}}$, and we know that

$$
c^{\prime \prime}=c c^{\prime} \prod_{\ell_{i} \in B} \ell_{i}^{-\alpha_{i}} \equiv 1 \quad(\bmod n)
$$

Because of the homomorphic property and the decryptions of the elements in $B$, we know how the decryption has been changed. And because $c^{\prime \prime}=\left[(1+n)^{m^{\prime}}\right]_{n^{2}}$ for some $m^{\prime}$, we know its decryption, and therefore the decryption of the original ciphertext $c$.

While this attack is of limited practical interest, it does provide some motivation to look at non-adaptive chosen ciphertext security for homomorphic cryptosystems.

### 5.2.1 Knowledge-of-exponent

Let $G$ be a group. For some groups, it seems difficult to find elements of the group, except by using the group operation to compute combinations of previously known elements. We formalise this in the following definition.

Definition. We say that a family of groups $\Gamma=\left\{\Gamma_{\tau}\right\}$ satisfies the knowledge-ofexponents assumption, if for any efficient algorithm $A$ that takes as input a group $G$ sampled from $\Gamma_{\tau}$ and a tuple of group elements $\left(g_{1}, \ldots, g_{m}\right)$, and outputs an element $x$ of $G$ with significant (in $\tau$ ) probability, there is a trace-variant $A^{\prime}$ of $A$ that also outputs a tuple of integers $\left(k_{1}, \ldots, k_{m}\right)$ such that $x=g_{1}^{k_{1}} \ldots g_{m}^{k_{m}}$.

This concept was first proposed by Damgård [12]. Other work on the subject is $[4,23]$. Given the subgroup membership problem $\mathcal{D D H}_{G^{\prime}}=\mathcal{S} \mathcal{M}_{(G, K)}$, they hypothesise that the subgroup $K$ has this property. That is, given a subgroup of $G^{\prime} \times G^{\prime}$ generated by $\left(g_{1}, g_{2}\right)$, the only way to find an element in this subgroup is to compute powers $\left(g_{1}^{k}, g_{2}^{k}\right)$, for various $k$.

We note that in the generic model described in Section 3.6.2, if the representation set $S$ is much larger than the group size, then the knowledge-of-exponents assumption holds.

We shall propose three new groups for which it is plausible that this property holds. The two first structures come from the subgroups of the symmetric subgroup membership problems $\mathcal{S O U _ { n }}$ and $\mathcal{G O} \mathcal{U}_{n}$. As for the Diffie-Hellman-based example, it simply seems difficult to generate elements in any other way than by computing powers of a given generator.

The third group structure that we shall propose is qualitatively different. Let $n=p q$ be an RSA modulus, and let $E$ be an elliptic curve defined over $\mathbb{Z}_{n}$, given by the equation

$$
y^{2} z \equiv x^{3}+a x z^{2}+b z^{3} \quad(\bmod n) .
$$

For an elliptic curve defined over a finite field $\mathbb{F}_{p}$, it is easy to find points on the curve. Simply choose an $x$-coordinate uniformly at random, and with probability approximately $1 / 2, x^{3}+a x+b$ will be a quadratic residue modulo $p$. Computing a square root $y$ of $x^{3}+a x+b$ modulo $p$ is easy, and $(x: y: 1)$ is an $\mathbb{F}_{p}$-rational point on $E$.

When we move to the ring $\mathbb{Z}_{n}$, the situation changes. Of course, we have the trivial point ( $0: 1: 0$ ), but it seems difficult to find any other point. Choosing an $x$-coordinate does not lead to a point, since computing square roots modulo $n$ is equivalent to factoring, and factoring seems hard. Starting with a $y$-coordinate, we get a cubic equation for the $x$-coordinate, but solving cubic equations seems to be hard.

If we are given a set of points, we can obviously use the group operation to find new points on the curve. We can also try to intersect the curve with other curves. One possibility is the quadratic curve

$$
C: y z \equiv a x^{2}+b x z+c z^{2} \quad(\bmod n)
$$

If we are given three distinct, non-trivial points $P_{1}, P_{2}, P_{3} \in E\left(\mathbb{Z}_{n}\right)$ that are not collinear, we can find such a quadratic curve through those points. Inserting the quadratic curve equation into the elliptic curve equation, we get a polynomial of degree 4 where we know three of the solutions, hence we can find the fourth and from that a point $P_{4}$ in $E\left(\mathbb{Z}_{n}\right)$.

To analyse this, we consider the situation modulo the prime factor $p$. By Bezout's theorem, $C$ and $E$ intersect in exactly 6 points when counting multiplicities. The four points $P_{1}, \ldots, P_{4}$ are given, and it is clear that a fifth intersection point is $O=(0: 1: 0)$. Since the line $z=0$ is tangent to both $E$ and $C$ in $O$, their intersection number at $O$ is greater than 1. This means that there are five points of intersection, and the divisor of $C$ is $\operatorname{div}(C)=\left(P_{1}\right)+\cdots+\left(P_{4}\right)+2(O)$.

If $C^{\prime}$ is the plane curve given by $z^{2}=0$, then $\operatorname{div}\left(C^{\prime}\right)=6(O)$. Let $f$ in the function field of $E$ be defined by $f(x, y, z)=\left(y z-a x^{2}-b x z-c z^{2}\right) / z^{2}$. It is then clear that

$$
\operatorname{div}(f)=\operatorname{div}(C)-\operatorname{div}\left(C^{\prime}\right)=0
$$

which gives us

$$
\left(P_{1}\right)-(O)+\left(P_{2}\right)-(O)+\left(P_{3}\right)-(O)+\left(P_{4}\right)-(O)=0
$$

or $P_{1}+P_{2}+P_{3}+P_{4}=O$. The four points sum to zero.
It seems as if most approaches to finding points in $E\left(\mathbb{Z}_{n}\right)$ either lead to polynomial equations modulo $n$, which seem to be hard to solve, or to trivial relations among points. It therefore seems plausible that the an elliptic curve modulo $n$ satisfies the knowledge-of-exponent assumption.

The qualitative difference between this group and the other group structures is that it is trivial to determine if a point $(x: y: z)$ is in $E\left(\mathbb{Z}_{n}\right)$ or not.

### 5.2.2 A security proof

We shall consider the cryptosystem $\Pi_{1}$ derived from a splitting problem $\mathcal{S} \mathcal{P}_{\Gamma}$, but let us suppose additionally that for every triple $(G, K, H) \in \Gamma_{\tau}, G$ is cyclic, and there is some group $\tilde{K}$ along with easily computable homomorphisms $\phi: G \rightarrow \tilde{K}$ and $\psi: \tilde{K} \rightarrow K$, such that $\phi \circ \psi$ is the map $x \mapsto x^{|H|}$, and $\phi$ restricted to $K$ is an isomorphism.

We shall assume that a description of $\tilde{K}$ is easy to recover from a description of $(G, K, H)$, and vice versa. We denote by $\tilde{\Gamma}$ the family of groups derived from $\Gamma$.

We shall also assume that a generator $g$ for $K$ is available for sampling elements. The public key for $\Pi_{1}$ is therefore ( $G, g, f$ ).

Theorem 5.5. Let $\Pi_{1}$ be derived from a splitting problem $\mathcal{S P}_{\Gamma}$ as described above. If the knowledge-of-exponent assumption holds for $\tilde{\Gamma}$, then $\Pi_{1}$ is weakly plaintext aware.

Proof. First of all, we note that if $p k=(G, g, f)$ is a valid public key for $\Pi_{1}$, then $p k^{\prime}=\left(G, g^{|H|}, f\right)$ is also a valid public key, and the two are interchangeable.

So assume that $B$ is a weak forger for $\Pi_{1}$, that is, $B$ takes as input a public key $p k$, and outputs a ciphertext $c \in G$. We use $B$ to construct an algorithm $A$ that takes as input a group $\tilde{K}$ along with a generator $\tilde{g}$, and outputs an element of $\tilde{K}$.

Given $\tilde{K}, A$ recovers $(G, K, H)$. It sets $g=\psi(\tilde{g})$, samples $f$ and runs $B$ on input $(G, g, f)$. $B$ outputs $c \in G$, and $A$ outputs $\phi(c)$.

Now it is clear that $A$ is an algorithm that takes as input the group $\tilde{K}$ and a generator for $\tilde{K}$, and outputs an element $\tilde{c}$ of $\tilde{K}$, so by the knowledge-ofexponent assumption, there exists some trace-variant $A^{\prime}$ that on the same input and random tape outputs the same element $\tilde{c}$ and an integer $k$ such that $\tilde{c}=\tilde{g}^{k}$.

Now it is quite clear that if $B$ is a forger, we can create a trace-variant $B^{\prime}$ of $B$ that outputs not only a ciphertext, but also the decryption of the ciphertext.
$B^{\prime}$ takes as input a public key $(G, g, f)$. It first runs $B$ on input $\left(G, g^{|H|}, f\right)$, and gets a ciphertext $c$ as output. Then it rewinds the random tape and runs $A^{\prime}$ on input of $(\tilde{K}, \phi(g))$. It is then quite clear that $A^{\prime}$, except with negligible probability, outputs $\tilde{c}$ and $k$ such that $\tilde{c}=\phi(c)$.

But then we have that

$$
\phi\left(c g^{-k|H|}\right)=\phi(c) \tilde{g}^{-k}=\tilde{c}\left(\tilde{c}^{-1}\right)=1
$$

which means that $c g^{-k|H|}$ must be in $H$, and since $g^{-k|H|}$ is in $K$, we can recover the decryption using $f^{-1}$. This also shows us how the weak knowledge extractor works.

It is now quite clear that by Theorem 3.4, we can conclude that $\Pi_{1}$ is secure against non-adaptive chosen ciphertext attacks, assuming that the derived subgroup membership problem is hard and that the knowledge-of-exponent assumption for $\tilde{\Gamma}$ holds. As a corollary, we get that $\Pi_{1}$ is secure against non-adaptive chosen ciphertext attacks in the generic model.

The two most interesting instantiations of this problem is for the subgroup variant of the Paillier cryptosystem, and the elliptic curve variant of the Paillier cryptosystem described by Galbraith [20].

### 5.3 A second homomorphic cryptosystem

Let $\Gamma_{\tau}$ be a set of tuples $(G, K, H, g)$ such that $\mathcal{S P}{ }_{(G, K, H)}$ is a splitting problem, $g$ is a generator for $K$, and $|H|$ has no small prime divisors. Let $\Gamma=\left\{\Gamma_{\tau}\right\}$. We shall be interested in the derived subgroup membership problem $\mathcal{S M}_{\Gamma}$, or the symmetric subgroup membership problem $\mathcal{S S} \mathcal{M}_{\Gamma}$.

The basic idea is that we shall do as for the ElGamal cryptosystem, but that the Diffie-Hellman key exchange part should be done in $K$ and the message should be either in $G$, or restricted to $H$. The exact construction of $\Pi_{2}$ from $\Gamma$ is described in Figure 5.3.

Proposition 5.6. The cryptosystem $\Pi_{2}$ is an homomorphic public key system.
Proof. Clear.
Theorem 5.7. The public key cryptosystem $\Pi_{2}$ is semantically secure for messages in $H$ if the subgroup membership problem $\mathcal{S M}_{\Gamma}$ is hard, and semantically secure for messages in $G$ if the symmetric subgroup membership problem $\mathcal{S S}_{\Gamma}$ is hard.

Proof. We assume that $A=\left(A_{1}, A_{2}\right)$ is a chosen plaintext adversary against indistinguishability for $\Pi_{2}$ with advantage $\epsilon$.

## Experiment 1.

Input: $(G, K, H, g), x \in G$.

1. $k \leftarrow\{0, \ldots,|G|-1\}, y \leftarrow g^{k}$.
2. $\left(m_{0}, m_{1}, o\right) \leftarrow A_{1}(G, K, H, g, k)$.
3. $b \leftarrow\{0,1\}$.
4. $e \leftarrow x^{k} m$.
5. $b^{\prime} \leftarrow A_{2}\left(m_{0}, m_{1}, x, e, o\right)$.
6. If $b=b^{\prime}$, output 1 , otherwise output 0 .

Output: 0 or 1 .

## Key generation.

Input: $\Gamma, \tau$.

1. $(G, K, H, g) \leftarrow \Gamma_{\tau}$.
2. $k \approx\{0, \ldots,|G|-1\}$.
3. $y \leftarrow g^{k}$.
4. $p k \leftarrow(G, K, H, g, y)$.
5. $s k \leftarrow(G, k)$.
6. Output ( $p k, s k$ ).

Output: A public-private key pair $(p k, s k)$.

## Encryption.

Input: $p k=(G, K, H, g, y), m \in G$.

1. $w \approx\{0, \ldots,|K|-1\}$.
2. $x \leftarrow g^{w}, z \leftarrow y^{w}, e \leftarrow z m$.
3. Output $(x, e)$.

Decryption.
Input: $s k=(G, k),(x, e) \in G \times G$.

1. $z \leftarrow x^{k}$.
2. $m \leftarrow e z^{-1}$.
3. Output $m$.

Output: A ciphertext $c \in G \times G$.
Output: A message $m \in G$.

Figure 5.3: The cryptosystem $\Pi_{2}$ derived from a (symmetric) subgroup membership problem. The message is restricted to $H$ if $\mathcal{S M}_{(G, H)}$ is not hard.

In a realisation of Experiment 1, we could perhaps not sample $k$ uniformly at random, but $\delta$-close to uniform (Lemma 2.3), where $\delta$ can be made arbitrarily small at little cost.

Note that Step 1 does exactly as $\mathcal{K}$ would do, and Step 4 does exactly as $\mathcal{E}$ would do, if it had first sampled $x$. (They would of course not know $k$, but would have sampled $x$ as $g^{w}$ and computed $y^{w}$.)

Let $T^{\prime}$ be the event that Experiment 1 outputs 1 when its input $x$ is in $K$. We note that if $x \in K$, then Experiment 1 proceeds exactly as a real attack, so

$$
\begin{equation*}
\operatorname{Adv}_{A}^{\mathrm{IND}} \leq\left|\operatorname{Pr}\left[T^{\prime}\right]-1 / 2\right| \tag{5.1}
\end{equation*}
$$

Let $T$ be the event that Experiment 1 outputs 1 when its input $x$ is in $G \backslash K$. We note that we can obviously derive an algorithm $A^{\prime}$ from Experiment 1 that distinguishes $K$ from $G \backslash K$ such that

$$
\begin{equation*}
\left|\operatorname{Pr}\left[T^{\prime}\right]-\operatorname{Pr}[T]\right| \leq \operatorname{Adv}_{A^{\prime}}^{\mathcal{S M}_{(G, K)}}+\delta \tag{5.2}
\end{equation*}
$$

Now we modify Experiment 1 as follows to get Experiment 1': We replace Step 4 with the following step:

$$
\text { 4. } z \leftarrow H, c \leftarrow x^{k} m z
$$

Let $T_{1}$ be the event that Experiment $1^{\prime}$ outputs 1 when its input $x$ is in $G \backslash K$.
By Lemma 4.12, the distributions of $\left(x, y, x^{k}\right)$ and $\left(x, y, x^{k} z\right)$ are $(|H|-$ $\phi(|H|)) /|H|$-close, and otherwise the two algorithms are identical, so

$$
\begin{equation*}
\left|\operatorname{Pr}[T]-\operatorname{Pr}\left[T_{1}\right]\right| \leq \frac{|H|-\phi(|H|)}{|H|} \tag{5.3}
\end{equation*}
$$

It is clear that if the message space is $H$, then the input to $A_{2}$ is independent of the bit $b$ chosen, so no matter what $A_{2}$ does, it guesses correctly with probability $1 / 2$, and

$$
\begin{equation*}
\operatorname{Pr}\left[T_{1}\right]=1 / 2 \tag{5.4}
\end{equation*}
$$

By combining (5.1)-(5.4), we immediately get that

$$
\operatorname{Adv}_{A}^{\mathrm{IND}} \leq \operatorname{Adv}_{A^{\prime}} \mathcal{S M}_{(G, K)}+\frac{|H|-\phi(|H|)}{|H|}+\delta
$$

Under the assumption that $\mathcal{S} \mathcal{M}_{(G, K)}$ is hard, this proves the first claim.
To prove the second claim, we first make a modification to Experiment $1^{\prime}$ to get Experiment $1^{\prime \prime}$ : We replace Step 4 with the following step:
4. $z \leftarrow G \backslash H, c \leftarrow x^{k} m z$.

Let $T_{2}$ be the event that Experiment $1^{\prime \prime}$ outputs 1 when its input $x$ is in $G \backslash K$. It is clear that if $z$ was sampled uniformly at random, then the input to $A_{2}$ would be independent of the bit $b$ chosen, so no matter what $A_{2}$ does, it would guess
correctly with probability $1 / 2$. By Lemma 2.2 , the uniform distribution on $G \backslash H$ is $2|H| /|G|$-close to the uniform distribution on $G$, so we have that

$$
\begin{equation*}
\left|\operatorname{Pr}\left[T_{2}\right]-1 / 2\right| \leq \frac{2|H|}{|G|} \tag{5.5}
\end{equation*}
$$

To bound the difference in probability between $T_{1}$ and $T_{2}$, we introduce the following experiment:

## Experiment 2.

Input: $(G, K, H, g), z \in G$.

1. $k \leftarrow\{0, \ldots,|G|-1\}, y \leftarrow g^{k}$.
2. $\left(m_{0}, m_{1}, o\right) \leftarrow A_{1}(G, K, H, g, k)$.
3. $b \leftarrow\{0,1\}$.
4. $x \leftarrow G \backslash K, e \leftarrow x^{k} m z$.
5. $b^{\prime} \leftarrow A_{2}\left(m_{0}, m_{1}, x, e, o\right)$.
6. If $b=b^{\prime}$, output 1 , otherwise output 0 .

Output: 0 or 1 .
Let $R^{\prime}$ be the event that Experiment 2 outputs 1 when its input $z$ is in $H$, and let $R$ be the corresponding event with the input $z$ is in $G \backslash H$.

It is quite clear that if the input $z$ to Experiment 2 is in $H$, then Experiment 2 behaves exactly as Experiment $1^{\prime}$ behaves when its input $x$ is in $G \backslash K$. Likewise, if $z \in G \backslash H$, then Experiment 2 behaves exactly as Experiment $1^{\prime \prime}$ behaves when its input $x$ is in $G \backslash K$. So we can derive an algorithm $A^{\prime \prime}$ from Experiment 2 such that

$$
\begin{equation*}
\left|\operatorname{Pr}\left[T_{1}\right]-\operatorname{Pr}\left[T_{2}\right]=\left|\operatorname{Pr}\left[R^{\prime}\right]-\operatorname{Pr}[R]\right| \leq \operatorname{Adv}_{A^{\prime \prime}}^{\mathcal{S}_{(G, H)}}+\delta\right. \tag{5.6}
\end{equation*}
$$

Combining (5.1)-(5.3), (5.5), and (5.6), we get that

$$
\operatorname{Adv}_{A}^{\mathrm{IND}} \leq \operatorname{Adv}_{A^{\prime}}^{\mathcal{S} \mathcal{M}_{(G, K)}}+\operatorname{Adv}_{A^{\prime \prime}}^{\mathcal{S}_{(G, K)}}+\frac{|H|-\phi(|H|)}{|H|}+2 \delta
$$

which proves the second claim.
Note that this cryptosystem does not require a trapdoor problem, unlike $\Pi_{1}$ and $\Pi_{1^{\prime}}$. We shall see more of the advantages of this approach in Section 7.

This variant of ElGamal can usefully be instantiated with the Decision Composite Residuosity problem (Section 4.5.4), to get an additively homomorphic cryptosystem. Various technical ${ }^{1}$ benefits of such an approach was discussed in [13].

It can also be usefully instantiated with the symmetric subgroup membership problems discussed in Section 4.5.6 and Section 4.5.7.

[^2]
## Chapter 6

## Key encapsulation methods

A key encapsulation method is a public key cryptosystem whose only job is to encapsulate keys for symmetric cryptosystems. The idea is that the weaker goal of transporting a random key could be easier to do than the more complicated task of transporting messages with an unknown distribution.

### 6.1 Security against passive attacks

Let $\mathcal{S P}_{(G, K, H)}$ be a splitting problem. The basic idea is that we can sample elements uniformly from $K$ and $H$ to be the random key, and hide them by multiplying. We recover the random key with the splitting algorithm. The key encapsulation method is described in Figure 6.1.

It is quite clear that the key encapsulation method $\Lambda_{1}$ is secure against key recovery if the splitting problem is hard. As we shall see, deciding if the correct key has been recovered is possible if either $K$ can be distinguished from $G \backslash K$ or if $H$ can be distinguished from $G \backslash H$.

Theorem 6.1. Let $\mathcal{T S P}_{\Gamma}$ be a trapdoor splitting problem, such that for every tuple $(G, K, H)$, the group order $|G|$ has no small prime factors. The key encapsulation method $\Lambda_{1}$ is semantically secure if and only if the symmetric subgroup membership problem $\mathcal{S S M}_{\Gamma}$ is hard.

Proof. Given a key-ciphertext pair $((x, y), c) \in G \times G \times G$, it is quite clear that if $x \in K$ or $y \in H$, then $(x, y)$ is with overwhelming probability the correct key.

Now we show that an algorithm for distinguishing the keys output by $\Lambda_{1}$ from random keys must lead to an algorithm for distinguishing either $K$ or $H$.

Suppose the sampling algorithm given by the symmetric subgroup membership problem samples elements $\delta$-close to the uniform distribution. We note that by using the sampling algorithms of the symmetric subgroup membership problem, and by Lemma 2.3, even if we do not know $|G|$, we can implement modified experiments whose underlying probability spaces are $O\left(\delta+\delta^{\prime}\right)$-close to the original experiments. The extra work needed is $O\left(\log 1 / \delta^{\prime}\right)$ exponentiations in $G$.

## Key generation.

Input: $\mathcal{T S P}_{\Gamma}, \tau$.

1. $(G, K, H, \sigma) \leftarrow \Gamma_{\tau}$.
2. $p k \leftarrow(G, K, H)$.
3. $s k \leftarrow(G, \sigma)$.
4. Output $(p k, s k)$.

Output: A public-private key pair $(p k, s k)$.

## Encryption.

Input: $p k=(G, K, H)$.

1. $x \approx K, y \approx H$.
2. Output $((x, y), x y)$.

Output: A key-ciphertext pair in $G \times G \times G$.

## Decryption.

Input: $s k=(G, \sigma), c \in G$.

1. $(x, y) \leftarrow \sigma(c)$.
2. Output $(x, y)$.

Output: A key in $G \times G$.

Figure 6.1: The key encapsulation $\operatorname{method} \Lambda_{1}$.

## Experiment 1.

Input: $(G, K, H), A, x \in G$.

1. $x^{\prime} \leftarrow K$.
2. $r \leftarrow\{0, \ldots,|G|-1\}$.
3. $y \leftarrow H$.
4. $b \leftarrow\{0,1\}$.
5. If $b=1$, then $\left(c, c^{\prime}, c^{\prime \prime}\right) \leftarrow\left(x^{r} x^{\prime}, y, x^{r} x^{\prime} y\right)$, otherwise $\left(c, c^{\prime}\right) \leftarrow G \times G$ and $c^{\prime \prime} \leftarrow c c^{\prime}$.
6. $b^{\prime} \leftarrow A\left(\left(c, c^{\prime}\right), c^{\prime \prime}\right)$.
7. If $b=b^{\prime}$, then output 1 , otherwise output 0 .

Output: 0 or 1.

## Experiment 2.

Input: $(G, K, H), A, y \in G$.

1. $x \leftarrow G$.
2. $y^{\prime} \leftarrow H$.
3. $r \leftarrow\{0, \ldots,|G|-1\}$.
4. $b \leftarrow\{0,1\}$.
5. If $b=1$, then $\left(c, c^{\prime}, c^{\prime \prime}\right) \leftarrow\left(x, y^{r} y^{\prime}, x y^{r} y^{\prime}\right)$, otherwise $\left(c, c^{\prime}\right) \leftarrow G \times G$ and $c^{\prime \prime} \leftarrow c c^{\prime}$.
6. $b^{\prime} \leftarrow A\left(\left(c, c^{\prime}\right), c^{\prime \prime}\right)$.
7. If $b=b^{\prime}$, then output 1 , otherwise output 0 .

Output: 0 or 1.

Figure 6.2: Experiments for the proof of Theorem 6.1.

Let $T_{1}$ be the event that Experiment 1 returns 1 when the input $x$ is in $K$. Let $T_{1}^{\prime}$ be the event that Experiment 1 returns 1 when the input $x$ is in $G \backslash K$. Let $T_{2}$ be the event that Experiment 2 returns 1 when the input $y$ is in $H$. Let $T_{2}^{\prime}$ be the event that Experiment 2 returns 1 when the input $y$ is in $G \backslash H$.

It is clear that if the input to Experiment 1 is in $K$, then the experiment proceeds exactly as an attack against the cryptosystem would do, except that in a real attack, $x$ would not be sampled uniformly but $\delta$-close to uniformly. In other words,

$$
\epsilon \leq\left|\operatorname{Pr}\left[T_{1}\right]-1 / 2\right|+\delta .
$$

Next, we see that if the input to Experiment 1 is in $G \backslash K$, then when $x$ is a generator for $G, x^{r}$ is uniformly distributed in $G$, so $\left(x^{r} x^{\prime}, y\right)$ is uniformly distributed in $G \times H$. By Lemma 2.5, the probability that $x$ is a generator is larger than $\phi(|G|) /|G|$.

If the input to Experiment 2 is in $H$, then no matter what $r$ is, $\left(x, y^{r} y^{\prime}\right)$ is uniformly distributed in $G \times H$. This means that

$$
\left|\operatorname{Pr}\left[T_{1}^{\prime}\right]-\operatorname{Pr}\left[T_{2}\right]\right| \leq \frac{|G|-\phi(|G|)}{|G|}
$$

Finally, suppose that the input $y$ to Experiment 2 is in $G \backslash H$, and that it is a generator for $G$. Then $\left(x, y^{r} y^{\prime}\right)$ is distributed uniformly in $G \times G$. It is then clear that

$$
\left|\operatorname{Pr}\left[T_{2}^{\prime}\right]-1 / 2\right| \leq \frac{|G|-\phi(|G|)}{|G|}
$$

Putting it all together, we get that

$$
\begin{aligned}
\epsilon \leq & \mid \operatorname{Pr}\left[T_{1}\right]-\operatorname{Pr}\left[T_{1}^{\prime}\right]+\operatorname{Pr}\left[T_{1}^{\prime}\right]-\operatorname{Pr}\left[T_{2}\right]+ \\
& \operatorname{Pr}\left[T_{2}\right]-\operatorname{Pr}\left[T_{2}^{\prime}\right]+\operatorname{Pr}\left[T_{2}^{\prime}\right]-1 / 2 \mid+\delta \\
\leq & \left|\operatorname{Pr}\left[T_{1}\right]-\operatorname{Pr}\left[T_{1}^{\prime}\right]\right|+\frac{|G|-\phi(|G|)}{|G|}+ \\
& \left|\operatorname{Pr}\left[T_{2}\right]-\operatorname{Pr}\left[T_{2}^{\prime}\right]\right|+\frac{|G|-\phi(|G|)}{|G|}+\delta .
\end{aligned}
$$

Since (up to $O\left(\delta^{\prime}\right)$ ) $\left|\operatorname{Pr}\left[T_{1}\right]-\operatorname{Pr}\left[T_{1}^{\prime}\right]\right|$ and $\left|\operatorname{Pr}\left[T_{2}\right]-\operatorname{Pr}\left[T_{2}^{\prime}\right]\right|$ are the advantages of some algorithms distinguishing $K$ from $G \backslash K$ and $H$ from $G \backslash H$, respectively, the assumption that the symmetric subgroup membership problem was hard means that $\epsilon$ must be negligible.

We can create a simple ElGamal-like cryptosystem using this key encapsulation method and the shift cipher in $G$. The symmetric cryptosystem has message space $G$, ciphertext space $G$ and key space $G \times G$. The encryption algorithm is $\mathcal{S E}((x, y), m)=x m$, and the decryption algorithm is $\mathcal{S D}((x, y), c)=c x^{-1}$.

When this cryptosystem is instantiated with the symmetric subgroup membership problem $\mathcal{S S M}_{(G, K, H)}=\mathcal{S O U}_{n}$ described in Section 4.5.6, we see that encryption requires one exponentiation in $K$ and one in $H$, while decryption requires essentially the same work.

Compared to ElGamal over $G$, which requires two exponentiations in $G$ for encryption and one for decryption, this is quite efficient. But $\Pi_{2}$ in Section 5.3 requires two exponentiations in $K$ for encryption, and only one exponentiation in $K$ for decryption. This is rather more efficient than the above system, and just as secure.

We conclude that the key encapsulation method $\Lambda_{1}$, while secure, is not very useful.

### 6.2 Security against active attacks

The key encapsulation method proposed in the previous section seems difficult to harden against adaptive attacks. We shall introduce a new assumption, and design a key encapsulation method secure against adaptive attacks in the random oracle model.

The basic idea is that if we pass the key in $\Lambda_{1}$ through a hash function and use the hash output as a key, it will be very difficult to recover the splitting from the key, and the decryption of other ciphertexts would not give any information about the challenge ciphertext. The reason this is not trivial, is that the adversary can use the decryption oracle to check if a group element is in one of the subgroups or not.

Definition. A gap splitting problem is a splitting problem $\mathcal{S P}_{(G, K, H)}$, but any algorithm that tries to solve the splitting problem has access to a splitting oracle that on input of $(x, y)$ and $c$ returns 1 if $(x, y)$ is the splitting of $c$, and 0 otherwise.

Gap problems were first defined by Okamoto and Pointcheval [32], and they use it to prove that their REACT conversion scheme [33] can be applied to e.g. ElGamal.

The gap problems are somewhat artificial, since the presence of an oracle is assumed. Therefore, any algorithm that solves the the gap problem is just a reduction from the splitting problem to the corresponding subgroup membership problems. If the subgroup membership problems are both easy, we get a real gap splitting problem.

The first real gap splitting problem appeared in [24]. The only known examples of such problems derive from certain elliptic curves with bilinear pairings. A bilinear pairing $e$ on a group $G$ is a map from $G \times G$ into a group $G^{\prime}$ such that for any $x \in G, e(x, x) \neq 1$, and for any $x, y \in G$ and $a, b \in \mathbb{Z}, e\left(x^{a}, y^{b}\right)=e(x, y)^{a b}$.

The most interesting variant is the one derived from the Computational DiffieHellman problem in $G$. It is clear that the pairing $e$ can give the solution to the Decision Diffie-Hellman. Given a tuple $\left(g, g^{x}, g^{y}, g^{z}\right)$, it is clear that $e\left(g^{x}, g^{y}\right)=$ $e(g, g)^{x y}$ and $e\left(g, g^{z}\right)=e(g, g)^{z}$. If $z \equiv x y(\bmod |G|)$, then the two pairings are equal, otherwise they are not.

We can also derive a real gap splitting problem from $\mathcal{S O U} \mathcal{E / F}_{p^{\prime}}$ discussed in Section 4.5.6, when the elliptic curve $E$ is chosen such that it is possible to compute a pairing on $E$. Suppose $\# E\left(\mathbb{F}_{p^{\prime}}\right)=n=p q$, and $P_{1}$ and $P_{2}$ are points of order $p$ and $q$, respectively. Then $P=P_{1}+P_{2}$ is a generator for $E\left(\mathbb{F}_{p^{\prime}}\right)$, and
for some $a, b \in \mathbb{Z}, P_{1}=a q P$ and $P_{2}=b p P$. It is then clear that we can use the pairing to decide subgroup membership. If a point $Q$ is in $\left\langle P_{1}\right\rangle$, then $Q=a^{\prime} q P$ for some $a^{\prime} \in \mathbb{Z}$, and we get that $e\left(Q, P_{2}\right)=e\left(\left(a^{\prime} q\right) P, b p P\right)=e(P, P)^{a^{\prime} b n}=1$.

We describe the key encapsulation method $\Lambda_{2}$ in Figure 6.3.

## Key generation.

Input: $\mathcal{T S P}_{\Gamma}, \tau$.

1. $(G, K, H, \sigma) \leftarrow \Gamma_{\tau}$.
2. Initialise random oracles $h_{1}: G \rightarrow\{0,1\}^{t}$ and $h_{2}:\{0,1\}^{2 t} \rightarrow\{0,1\}^{t}$.
3. $p k \leftarrow\left(G, K, H, h_{1}, h_{2}\right)$.
4. $s k \leftarrow\left(G, \sigma, h_{1}, h_{2}\right)$.
5. Output ( $p k, s k$ ).

Output: A public-private key pair $(p k, s k)$.

## Encryption.

Input: $p k=\left(G, K, H, h_{1}, h_{2}\right)$.

1. $x \approx K, y \approx H$.
2. Output $\left(h_{2}\left(h_{1}(x), h_{1}(y)\right), x y\right)$.

Output: A key-ciphertext pair in $\{0,1\}^{t} \times G$.

## Decryption.

Input: $s k=\left(G, \sigma, h_{1}, h_{2}\right), c \in G$.

1. $(x, y) \leftarrow \sigma(c)$.
2. Output $h_{2}\left(h_{1}(x), h_{1}(y)\right)$.

Output: A key in $\{0,1\}^{t}$.

Figure 6.3: The key encapsulation method $\Lambda_{2}$.
It is quite obvious that if the underlying splitting problem is hard, this KEM is secure against passive attacks, because the adversary has to compute the splitting of the ciphertext to get any information about the resulting hash value.

Theorem 6.2. The key encapsulation method $\Lambda_{2}$ is secure against adaptive attacks in the random oracle model if and only if the gap splitting problem is hard.

Proof. We are given an active adversary against the KEM, and also an oracle for deciding if a splitting is correct or not. First, we assume that we actually know the splitting algorithm $\sigma$, and we design an algorithm $D$ that shall play the role of the decryption oracle and the random oracles.

Then we create a modified algorithm $D^{\prime}$ that answers the same queries, but without recourse to the splitting algorithm $\sigma$, and such that this algorithm looks like $D$ from the adversary's point of view.
$D$ answers queries as follows. Queries to $h_{1}$ and $h_{2}$ are handled exactly as a real random oracle would handle them (see Section 3.6.1), using a list $L$ of query-answer pairs for the oracle $h_{1}$ and a list $L^{\prime}$ of query-answer pairs for $h_{2}$. Decryption oracle queries are handled by using the $\sigma$ algorithm to split the ciphertext, then by querying itself for the relevant hash values.

This completes the description of the original algorithm, and it is clear that it will work exactly as the real oracles would work.

Now we describe the modified algorithm $D^{\prime}$. The list $L$ is split into three lists, $L_{G}, L_{K}$ and $L_{H}$. Whenever a $h_{1}$-query $x$ is received that has not previously been answered, the splitting oracle is used to check if $x$ is in $K$ or $H$. If it is in $K$, the response is recorded in $L_{K}$, likewise for $H$, and if it is not in $K \cup H$, then it is recorded in $L_{G}$.

The list $L^{\prime}$ remains the same.
We also need a third list $L^{\prime \prime}$ containing decryption oracle query-answer pairs. Whenever a new query $c$ is received by the decryption oracle, it first checks to see if the lists $L_{K}$ and $L_{H}$ contain queries $(x, b)$ or $\left(y, b^{\prime}\right)$ such that $\left(x, c x^{-1}\right)$ or $\left(c y^{-1}, y\right)$ is a splitting of $c$. This can be done using the splitting oracle.

If there are two queries $(x, b)$ and $\left(y, b^{\prime}\right)$, the algorithm makes the $h_{2}$-query $\left(b, b^{\prime}\right)$ to itself, receiving an answer $b^{\prime \prime}$. It adds the query $\left(c, b^{\prime \prime}\right)$ to $L^{\prime \prime}$, and outputs $b^{\prime \prime}$.

If there is only one query, say $(x, b) \in L_{K}$, but no query $\left(y, b^{\prime}\right)$ in $L_{H}, b^{\prime}$ is sampled uniformly at random, the query $\left(c x^{-1}, b^{\prime}\right)$ is inserted into $L_{H}$, and the algorithm makes the $h_{2}$-query ( $b, b^{\prime}$ ) to itself, receiving an answer $b^{\prime \prime}$. It then adds the query $\left(c, b^{\prime \prime}\right)$ to $L^{\prime \prime}$, and outputs $b^{\prime \prime}$.

If there are no such queries, $b^{\prime \prime}$ is sampled uniformly at random the query $\left(c, b^{\prime \prime}\right)$ is added to $L^{\prime \prime}$, and $b^{\prime \prime}$ is output.

To preserve consistency, we now have to check when receiving a new $h_{1}$-query $x \in K$ if for any $\left(c, b^{\prime \prime}\right) \in L^{\prime \prime},\left(x, c x^{-1}\right)$ is a splitting. If not, we proceed normally, otherwise, we need to sample $b$ uniformly at random and $b^{\prime}$ uniformly at random, subject to the restriction that there be no query $\left(b, b^{\prime}, b^{\prime \prime \prime}\right) \in L^{\prime}$, for some $b^{\prime \prime \prime} \neq b^{\prime \prime}$. Then we add $(x, b)$ to $L_{K},\left(c x^{-1}, b^{\prime}\right)$ to $L_{H},\left(b, b^{\prime}, b^{\prime \prime}\right)$ to $L^{\prime}$.

It is quite clear that the algorithm $D^{\prime}$ behaves consistently and is indistinguishable from $D$. Since it does not use the splitting algorithm $\sigma$, it is also clear that the adversary can run $D^{\prime}$ for himself. The end result is that the decryption oracle does not give the adversary anything he could not get from the splitting oracle.

Now we prove the reverse direction, by showing how we can use the decryption oracle to get a splitting oracle that is correct except with negligible error. So suppose we receive a query $c$ and a potential splitting $(x, y)$. Obviously, $c$ is a valid ciphertext. If $(x, y)$ really is the splitting, then $h_{2}\left(h_{1}(x), h_{1}(y)\right)$ is the correct decryption of $c$. Otherwise, if $\left(x^{\prime}, y^{\prime}\right)$ is the splitting of $c,\left(h_{1}(x), h_{1}(y)\right)$ collides with $\left(h_{1}\left(x^{\prime}\right), h_{1}\left(y^{\prime}\right)\right)$ with probability $\approx 2^{-2 t}$. If it is distinct, the output of $h_{2}$ collides with probability $\approx 2^{-t}$. In other words, our simulated splitting oracle never answers "no" incorrectly, and answers "yes" incorrectly with negligible probability.

We briefly consider the performance of this KEM, compared to Diffie-Hellman key encapsulation method (DH-KEM) [1, 16]. It is possible to prove DH-KEM secure in the random oracle model under a gap Diffie-Hellman problem, much as in the above proof.

Note that it would be possible to do DH-KEM in one of the two subgroups. The cost for DH-KEM is two subgroup exponentiations for encryption and one for decryption.

Encryption with $\Lambda_{2}$ costs the same as sampling one random element from each subgroup. Typically, this is done by exponentiation, but in some cases there are faster methods available. Decryption is done using one exponentiation in the group. This will typically cost as much as one exponentiation in each subgroup.

So we see that we can gain an advantage in encryption speed if sampling random elements in at least one group is faster than exponentiations in the subgroups.

When $\Lambda_{2}$ is instantiated with a Computational Diffie-Hellman problem $\mathcal{C D H}_{G^{\prime}}$, where it is possible to quickly sample elements uniformly at random from $G^{\prime}$, the resulting scheme is is exactly as fast as DH-KEM in the subgroup $H$ (which is isomorphic to $G^{\prime}$ ). The drawback is that the ciphertexts are twice as long.

We can consider DH-KEM to be a special case of our construction. First we note that if $(x, y)$ and $\left(x^{\prime}, y^{\prime}\right)$ are in $G=G^{\prime} \times G^{\prime}$, it is clear that they have the same projection on $K$ if and only if $x=x^{\prime}$. We consider two ciphertexts to be equivalent if their first coordinates are the equal. Then we change the hash function $h_{2}$ so that it simply outputs the $t$ first bits of its input. Finally, we change the decryption oracle, so that it refuses to decrypt ciphertexts that are equivalent to the challenge ciphertext.

It is now clear that the second coordinate of the ciphertext can be removed, and the resulting scheme is exactly $\mathrm{DH}-\mathrm{KEM}$.

When $\Lambda_{2}$ is instantiated with the Computational Composite Residuosity problem $\mathcal{C C R}_{n}$, sampling elements in $H$ is essentially free, so encryption costs one exponentiation in $K$. Decryption can be done using one exponentiation in $K$ and one in $H$ (which is essentially free). This instantiation is therefore much faster at encryption than DH-KEM in $K$, and marginally slower at decryption.

However, it is clear that DH-KEM in $K$ is exactly as hard as DH-KEM in a subgroup of $\mathbb{Z}_{n}^{*}$, and then DH-KEM becomes about twice as fast as $\Lambda_{2}$ for encryption, and four times as fast for decryption, because the numbers it works with are half the size. We also recall that $\mathcal{F A C}_{n}$ cannot be hard unless $\mathcal{C D} \mathcal{H}_{\mathbb{Z}_{n}^{*}}$ is hard, and $\mathcal{C C} \mathcal{R}_{n}$ cannot be hard unless $\mathcal{F A C}_{n}$ is hard. Therefore, DH-KEM in $\mathbb{Z}_{n}^{*}$ is faster than, and most likely not less secure than, $\Lambda_{2}$ instantiated with $\mathcal{C C R}{ }_{n}$.

The final interesting case is to instantiate $\Lambda_{2}$ with the problem $\mathcal{G O} \mathcal{U}_{n}^{\prime}$ described in Section 4.5.7. Note that in this case, it would seem safe to make $H$ fairly small, say about one quarter of the size of $K$. Exponentiations in $H$ are therefore
much faster than in $K$, and splitting can be done with one exponentiation in $H$ and one in $K$. This means that $\Lambda_{2}$ will be quite a lot faster at encryption than DH-KEM in $K$, and it will be slightly slower at decryption.

If we compare with DH-KEM in a subgroup $G^{\prime}$ of a finite field, however, it seems clear that the subgroup $K$ has to be twice as big as $G^{\prime}$, and therefore that exponentiations in $K$ take double the time required for exponentiations in $G^{\prime}$. One may make the assumption that the splitting problem remains hard even if exponents are restricted to a much smaller size, but that seems to be stretching the security assumptions for $\mathcal{G O} \mathcal{U}_{n}^{\prime}$.

Our conclusion therefore seems to be that, even though we can find underlying group structures where $\Lambda_{2}$ will be faster than DH-KEM, there are group structures available for DH -KEM that are not available for $\Lambda_{2}$, and they make DH-KEM a better choice.

Just as for DH-KEM [1], it is possible to formulate a decision problem based on the gap splitting problem and a hash function.

Definition. Let $(G, K, H, \sigma)$ be a trapdoor splitting problem, and let $h: G \times G \rightarrow$ $\{0,1\}^{t}$ be a hash function.

The hash splitting problem is the following problem: The instances are pairs $(c, b) \in G \times\{0,1\}^{t}$, and the answer is 1 if $h(\sigma(c))=b, 0$ otherwise. The instances are sampled by sampling $c$ uniformly at random, then letting $b=h(c)$ with probability 1 , otherwise sampling $b$ uniformly at random.

The oracle splitting problem is the same problem, but any adversary against the problem gets access to a restricted oracle that computes $h$ for all group elements except the challenge element.

It is quite clear that $\Lambda_{2}$ is secure against active attacks if the oracle splitting problem is hard for $h$. Indeed, it is just a reformulation of what must be true for the key encapsulation mechanism to be secure, and we may as well assume that our particular instantiation of $\Lambda_{2}$ is secure.

## Chapter 7

## A secure cryptosystem

The goal of this chapter is to describe a practical public key cryptosystem that is secure against adaptive attacks in the standard model. To get the cryptosystem, we shall extend the construction of Cramer and Shoup [10] slightly.

### 7.1 Hash proof systems

The basic idea of Cramer and Shoup [10] is that if it is possible to define a function on some set $G$, but only give away a description of the function on a subset $K$, this is a useful primitive for designing cryptosystems. Our discussion follows [10] closely, even though it is superficially different.

Let $G$ be a set, and let $K$ be a subset of $G$. We say that $W$ is a witness set for $K$ if there is an easily computable bijection $\rho: W \rightarrow K$. This bijection allows one to prove that an element $x \in G$ really is in $K$ by presenting an element $w \in W$ such that $\rho(w)=x$. (This obviously assumes that it is easy to recognise elements of $W$.)

For arbitrary sets $S, S^{\prime}$, denote by $\operatorname{Map}\left(S, S^{\prime}\right)$ the set of maps from $S$ to $S^{\prime}$. Let $L$ be a group. We are interested in looking at maps from $G$ to $L$. There is a natural map $\operatorname{Map}(G, L) \rightarrow \operatorname{Map}(K, L)$ given by restriction. The bijection $\rho$ gives us a bijection $\rho^{*}: \operatorname{Map}(K, L) \rightarrow \operatorname{Map}(W, L)$. We also denote the natural map $\operatorname{Map}(G, L) \rightarrow \operatorname{Map}(W, L)$ as $\rho^{*}$. It is clear that the following diagram commutes:


A projective hash family is a tuple $(G, K, L, W, \rho, M)$, where $G$ is a set, $K$ is a subset of $G, L$ is a group, $W$ is a witness set for $K$ with isomorphism $\rho, M$ is a subset of $\operatorname{Map}(G, K)$, and for any $f \in M$, the image of $K$ under $f$ is a subgroup of $L$. We also suppose that $L$ has a subgroup $L^{\prime}$, such that $L^{\prime} \cap f(K)=\{1\}$ and
$L=L^{\prime} f(K)$. This gives us a subgroup membership problem $\mathcal{S M}_{\left(L, L^{\prime}\right)}$. We say that $\mathcal{S M}_{\left(L, L^{\prime}\right)}$ is the subgroup membership problem associated to $(G, K)$ by the projective hash family.

Let $(G, K, L, W, \rho, M)$ be a projective hash family, and let $F$ be uniformly distributed in $M$. The projective hash family is $\epsilon$-universal if for any $f \in \rho^{*}(M)$, $x \in G \backslash K$ and $y \in L$, we have that

$$
\operatorname{Pr}\left[F(x)=y \mid \rho^{*}(F)=f\right] \leq \epsilon
$$

The projective hash family is $\epsilon$-universal-2 if for any $f \in \rho^{*}(M), x_{0} \in G \backslash K$, $x \in G \backslash\left(K \cup\left\{x_{0}\right\}\right)$ and $y, y_{0} \in L$, we have that

$$
\begin{equation*}
\operatorname{Pr}\left[F(x)=y \mid F\left(x_{0}\right)=y_{0} \wedge \rho^{*}(F)=f\right] \leq \epsilon \tag{7.1}
\end{equation*}
$$

It is clear that $\epsilon$-universal follows from $\epsilon$-universal-2.
Let $(G, K, L, W, \rho, M)$ be a projective hash family. Let $F$ be uniformly distributed on $M$, let $X$ be uniformly distributed in $G \backslash K$, and let $Y$ be uniformly distributed in $L^{\prime}$. Define the two random variables $U=\left(X, \rho^{*}(F), F(X)\right)$ and $V=\left(X, \rho^{*}(F), F(X) Y\right)$. We say that the projective hash family is $\epsilon$-smooth if

$$
\operatorname{Dist}(U, V) \leq \epsilon
$$

Let $\mathcal{S} \mathcal{M}_{\Gamma}$ be a subgroup membership problem. A hash proof system $P$ for $\mathcal{S} \mathcal{M}_{\Gamma}$ is a function that to each pair $(G, K) \in \Gamma$ assigns a tuple $(L, W, \rho, M)$ such that $(G, K, L, W, \rho, M)$ is a projective hash family.

An extended hash proof system $\hat{P}$ for $\mathcal{S M}_{\Gamma}$ is a function that to each pair $(G, K) \in \Gamma$ assigns a tuple $(S, \hat{L}, W, \hat{\rho}, \hat{M})$ such that $(G \times S, K \times S, \hat{L}, W, \hat{\rho}, \hat{M})$, where $S$ is some set depending on $G$, is a projective hash family.

A hash proof system $P$ is $\epsilon(\tau)$-smooth if the projective hash families derived using $P$ are $\epsilon(\tau)$-smooth except with negligible (in $\tau$ ) probability.

An extended hash proof system $\hat{P}$ is $\epsilon(\tau)$-universal- 2 if the projective hash families derived using $\hat{P}$ are $\epsilon(\tau)$-universal-2 except with negligible (in $\tau$ ) probability.

A hash proof system also has to provide algorithms to sample the sets $W$, $M$ and $\hat{M}$. These algorithms should sample the sets $\delta$-close to uniform, and we denote sampling using these algorithms by $w \approx W, f \approx M$, and $\hat{f} \approx \hat{\mathcal{M}}$.

### 7.2 The general construction

Suppose we have a subgroup membership problem $\mathcal{S}_{\Gamma}$, a hash proof system $P$ for $\mathcal{S} \mathcal{M}_{\Gamma}$, and an extended hash proof system $\hat{P}$ for $\mathcal{S} \mathcal{M}_{\Gamma}$ such that if $(L, W, \rho, M)$ is the tuple assigned to $(G, K)$ by $P$, then the tuple assigned to $(G, K)$ by $\hat{P}$ is $(L, \hat{L}, W, \hat{\rho}, \hat{M})$.

We can derive a cryptosystem from the two hash proof systems. It is described in Figure 7.1.

## Key generation.

Input: $\mathcal{S} \mathcal{M}_{\Gamma}, P, \hat{P}, \tau$.

1. $(G, K) \leftarrow \Gamma_{\tau}$.
2. $(L, W, \rho, M) \leftarrow P(G, K)$.
3. $(L, \hat{L}, W, \hat{\rho}, \hat{M}) \leftarrow \hat{P}(G, K)$.
4. $f \approx M$.
5. $\hat{f} \approx \hat{M}$.
6. $s k \leftarrow(G, L, \hat{L}, f, \hat{f})$.
7. $p k \leftarrow\left(G, W, L, \hat{L}, \rho, \rho^{*}(f), \hat{\rho}^{*}(\hat{f})\right)$.

Output: $(p k, s k)$.

## Encryption.

## Decryption.

Input: $p k=\left(G, W, L, \hat{L}, \rho^{*}(f), \hat{\rho}^{*}(\hat{f})\right)$, Input: $s k=(G, L, \hat{L}, f, \hat{f}),(x, e, \hat{y})$. $m \in L$.

1. $w \approx W$.
2. $\hat{y}^{\prime} \leftarrow \hat{f}(x, e)$.
3. $x \leftarrow \rho(w)$.
4. $y \leftarrow \rho^{*}(f)(w)$.
5. $e \leftarrow y m$.
6. $\hat{y} \leftarrow \rho^{*}(\hat{f})(w, e)$.
7. If $\hat{y}^{\prime} \neq \hat{y}$, output $\perp$.
8. $y \leftarrow f(x)$.
9. $m \leftarrow e y^{-1}$.
10. Output $m$.

Output: $(x, e, \hat{y})$.

Figure 7.1: The cryptosystem $\Pi_{3}$ derived from the hash proof systems $P$ and $\hat{P}$, and the subgroup membership problem $\mathcal{S M}_{\Gamma}$.

The security analysis closely follows the analysis in [10].
Suppose that $P$ is $\epsilon$-smooth, that $\hat{P}$ is $\epsilon^{\prime}$-universal-2, that the sampling algorithms for $P$ and $\hat{P}$ are $\delta$-close to uniform, and that the sampling algorithms for the subgroup membership problem are $\delta^{\prime}$-close to uniform.

Suppose $A=\left(A_{1}, A_{2}\right)$ is a chosen ciphertext adversary against $\Pi_{3}$. We shall use the following experiment to construct a distinguisher $A^{\prime}$ for $(G, K)$.

## Experiment 1.

Input: $A=\left(A_{1}, A_{2}\right),(G, K), P, \hat{P}, x_{0} \in G$.

1. $(L, W, \rho, M) \leftarrow P(G, K)$.
2. $(L, \hat{L}, W, \hat{\rho}, \hat{M}) \leftarrow \hat{P}(G, K)$.
3. $f \approx M, \hat{f} \approx \hat{M}$.
4. $s k \leftarrow(G, L, \hat{L}, f, \hat{f})$.
5. $p k \leftarrow\left(G, W, L, \hat{L}, \rho, \rho^{*}(f), \hat{\rho}^{*}(\hat{f})\right)$.
6. Initialise decryption oracle $\mathcal{D}_{s k}$.
7. $\left(m_{0}, m_{1}, s\right) \leftarrow A_{1}(p k)$, giving $A_{1}$ access to $\mathcal{D}_{s k}$.
8. $b \leftarrow\{0,1\}$.
9. $y_{0} \leftarrow f\left(x_{0}\right), e_{0} \leftarrow y_{0} m_{b}, \hat{y}_{0} \leftarrow \hat{f}\left(x_{0}, e_{0}\right)$.
10. Initialise restricted decryption oracle $\mathcal{D}_{s k}^{\prime}$.
11. $b^{\prime} \leftarrow A_{2}\left(p k, m_{0}, m_{1}, s, x_{0}, e_{0}, \hat{y}_{0}\right)$, giving $A_{2}$ access to $\mathcal{D}_{s k}^{\prime}$.
12. If $b=b^{\prime}$, output 1 , otherwise output 0 .

Output: 0 or 1 .
Note that Steps 1-5 do exactly as the key generation algorithm would do.
Let $T^{\prime}$ be the event that Experiment 1 outputs 1 when the input $x_{0}$ is in $K$. Since Step 9 produces exactly the same result as the encryption algorithm when the input $x_{0} \in K$, it is clear that the only difference between Experiment 1 and a real attack is that $x_{0}$ has been sampled uniformly from $K$, and not via the sampling algorithm for $W$. Since Experiment 1 outputs 1 when the adversary wins, we have that

$$
\begin{equation*}
\operatorname{Adv}_{A}^{C C A 2} \leq\left|\operatorname{Pr}\left[T^{\prime}\right]-1 / 2\right|+\delta \tag{7.2}
\end{equation*}
$$

since the sampling algorithm for $W$ is $\delta$-close to uniform.
Let $T$ be the event that Experiment 1 outputs 1 when the input $x_{0}$ is in $G \backslash K$. It is quite clear that from Experiment 1 we can derive an algorithm $A^{\prime}$ for distinguishing $K$ from $G \backslash K$ such that

$$
\begin{equation*}
\left|\operatorname{Pr}\left[T^{\prime}\right]-\operatorname{Pr}[T]\right| \leq \operatorname{Adv}_{A^{\prime}}^{\mathcal{S} \mathcal{M}_{(G, K)}} \tag{7.3}
\end{equation*}
$$

To analyse the event $T$, we shall make a series of modifications to Experiment 1. We number the modified experiments as $1^{\prime}, 1^{\prime \prime}$, etc. Note that these modifications need not be efficiently implementable.

First modification We change Step 3 so that $f$ and $\hat{f}$ are sampled from the uniform distribution, and not using the algorithms provided by the hash proof systems.

Let $T_{1}$ be the event that Experiment $1^{\prime}$ outputs 1 when the input $x_{0}$ is in $G \backslash K$. Since the algorithms provided by the hash proof systems were $\delta$-close to uniform, we obviously have that

$$
\begin{equation*}
\left|\operatorname{Pr}[T]-\operatorname{Pr}\left[T_{1}\right]\right| \leq 2 \delta . \tag{7.4}
\end{equation*}
$$

Second modification We change the decryption oracles so that they refuse to decrypt a ciphertext $(x, e, \hat{y})$ if $x \notin K$. Let $T_{2}$ be the event that Experiment $1^{\prime \prime}$ outputs 1 when the input $x_{0}$ is in $G \backslash K$.

It is clear that this modification only affects the outcome if the adversary produces a valid ciphertext $\left(x^{\prime}, e^{\prime}, \hat{y}^{\prime}\right)$ with $x \notin K$, so $\left|\operatorname{Pr}\left[T_{2}\right]-\operatorname{Pr}\left[T_{1}\right]\right|$ is upperbounded by the probability of this happening.

We first consider the decryption queries made by $A_{1}$. We condition on a fixed input $\rho^{*}(f)$ and $\hat{\rho}^{*}(\hat{f})$, and a fixed random tape. Note that this completely determines the output of $A_{1}$. If $A_{1}$ makes any decryption query $(x, e, \hat{y})$, with $x \notin K$, then since $\hat{P}$ is $\epsilon^{\prime}$-universal, we have that $(x, e, \hat{y})$ is valid with probability at most $\epsilon^{\prime}$.

Next we consider the decryption queries made by $A_{2}$. We additionally condition on a fixed value of $b$, the input $x_{0}$, the $\hat{y}_{0}$ computed in Step 9 , along with the random tape of $A_{2}$. If $A_{2}$ makes any decryption query ( $x, e, \hat{y}$ ), with $x \notin K$, we have two possibilities. If $(x, e)=\left(x_{0}, e_{0}\right)$, then we must have $\hat{y} \neq \hat{y}_{0}$, hence the ciphertext must be invalid. If $(x, e) \neq\left(x_{0}, e_{0}\right)$, then since $\hat{P}$ is $\epsilon^{\prime}$-universal-2, we have that $(x, e, \hat{y})$ is valid with probability at most $\epsilon^{\prime}$.

If $A_{1}$ and $A_{2}$ make $Q$ decryption queries in total, we get that

$$
\begin{equation*}
\left|\operatorname{Pr}\left[T_{2}\right]-\operatorname{Pr}\left[T_{1}\right]\right| \leq Q \epsilon^{\prime} \tag{7.5}
\end{equation*}
$$

Third modification We change Step 9 to be
9. $y^{\prime} \leftarrow L^{\prime}, y_{0} \leftarrow f\left(x_{0}\right), e_{0} \leftarrow y_{0} m_{b} y^{\prime}, \hat{y}_{0} \leftarrow \hat{f}\left(x_{0}, e_{0}\right)$.

Let $T_{3}$ be the event that Experiment $1^{\prime \prime \prime}$ outputs 1 when the input $x_{0}$ is in $G \backslash K$.
Since $A_{1}$ and $A_{2}$ cannot query the decryption oracle with ciphertexts ( $x, e, \hat{y}$ ) where $x \notin K$, their only information about $f$ is $\rho^{*}(f)$. Since $P$ is $\epsilon$-smooth, we get that

$$
\begin{equation*}
\left|\operatorname{Pr}\left[T_{3}\right]-\operatorname{Pr}\left[T_{2}\right]\right| \leq \epsilon \tag{7.6}
\end{equation*}
$$

Fourth modification We change Step 9 to be
9. $y^{\prime} \leftarrow L \backslash L^{\prime}, y_{0} \leftarrow f\left(x_{0}\right), e_{0} \leftarrow y_{0} m_{b} y^{\prime}, \hat{y}_{0} \leftarrow \hat{f}\left(x_{0}, e_{0}\right)$.

Let $T_{4}$ be the event that Experiment $1^{\prime \prime \prime \prime}$ outputs 1 when the input $x_{0}$ is in $G \backslash K$.
It is quite clear that if $y^{\prime}$ had been sampled uniformly from $L$, then there would be no information about $m_{b}$ present in the ciphertext, and the probability
that Experiment $1^{\prime \prime \prime \prime}$ output 1 when the input $x_{0}$ was in $G \backslash K$ would be $1 / 2$. Since Experiment $1^{\prime \prime \prime \prime}$ samples from $L \backslash L^{\prime}$, we get that

$$
\begin{equation*}
\left|\operatorname{Pr}\left[T_{4}\right]-1 / 2\right| \leq \frac{2\left|L^{\prime}\right|}{|L|} \tag{7.7}
\end{equation*}
$$

by Lemma 2.2.
We need to bound $\mid \operatorname{Pr}\left[T_{4}\right]-\operatorname{Pr}\left[T_{3} \mid\right.$. To do this, we introduce another experiment.

## Experiment 2.

Input: $A=\left(A_{1}, A_{2}\right),(G, K), P, \hat{P}, y^{\prime} \in L$.
Steps $1-8$ are as in Experiment 1.
9. $x_{0} \leftarrow G \backslash K, y_{0} \leftarrow f\left(x_{0}\right), e_{0} \leftarrow y_{0} m_{b} y^{\prime}, \hat{y}_{0} \leftarrow \hat{f}\left(x_{0}, e_{0}\right)$.

Steps 10-12 are as in Experiment 1.
Output: 0 or 1.
It is quite clear that we can repeat the two first modifications to Experiment 1 on Experiment 2, and the analysis remains the same. Let $R^{\prime}$ be the event that Experiment $2^{\prime \prime}$ outputs 1 when the input $y^{\prime}$ is in $L^{\prime}$, and let $R$ be the event that Experiment $2^{\prime \prime}$ outputs 1 when the input $y^{\prime}$ is in $L \backslash L^{\prime}$.

If the input $y^{\prime}$ to Experiment $2^{\prime \prime}$ is in $L^{\prime}$, then it is quite clear that it behaves exactly as Experiment $1^{\prime \prime \prime}$. Hence, $\operatorname{Pr}\left[R^{\prime}\right]=\operatorname{Pr}\left[T_{3}\right]$.

If the input $y^{\prime}$ to Experiment $2^{\prime \prime}$ is in $L \backslash L^{\prime}$, then it is quite clear that it behaves exactly as Experiment $1^{\prime \prime \prime \prime}$. Hence, $\operatorname{Pr}[R]=\operatorname{Pr}\left[T_{4}\right]$.

It is also quite clear that we from Experiment 2 can derive an algorithm $A^{\prime \prime}$ to distinguish $L^{\prime}$ from $L \backslash L^{\prime}$, by sampling $x_{0}$ not uniformly from $G \backslash K$, but via the subgroup membership problem's algorithms, and that

$$
\begin{equation*}
\left|\operatorname{Pr}\left[T_{4}\right]-\operatorname{Pr}\left[T_{3}\right]\right|=\left|\operatorname{Pr}[R]-\operatorname{Pr}\left[R^{\prime}\right]\right| \leq \operatorname{Adv}_{A^{\prime \prime}}^{\mathcal{S} \mathcal{M}_{\left(L, L^{\prime}\right)}}+2 \delta+\delta^{\prime}+Q \epsilon^{\prime} . \tag{7.8}
\end{equation*}
$$

Summing up Putting together (7.2)-(7.8), it is clear that we have proved the following theorem.

Theorem 7.1. Let $\Pi_{3}$ be the cryptosystem described in Figure 7.1, based on a subgroup membership problem $\mathcal{S}_{(G, K)}$ and hash proof systems $P$ and $\hat{P}$. Let $L$ be the group associated to $G$ by $P$, and let $L^{\prime}$ be the subgroup of $L$. Suppose that $P$ is $\epsilon$-smooth, that $\hat{P}$ is $\epsilon^{\prime}$-universal-2, that the sampling algorithms for $P$ and $\hat{P}$ are $\delta$-close to uniform, and that the sampling algorithms for the subgroup membership problem are $\delta^{\prime}$-close to uniform. Then for any chosen ciphertext adversary $A$ against $\Pi_{3}$, we have that

$$
\operatorname{Adv}_{A}^{C C A 2}(\tau) \leq \operatorname{Adv}_{A^{\prime}}^{\mathcal{S} \mathcal{M}_{(G, K)}}+\operatorname{Adv}_{A^{\prime \prime}}^{\mathcal{S} \mathcal{M}_{\left(L, L^{\prime}\right)}}+4 \delta+\delta^{\prime}+2 Q \epsilon^{\prime}+\epsilon+\frac{2\left|L^{\prime}\right|}{|L|}
$$

where $A^{\prime}$ and $A^{\prime \prime}$ are algorithms that invoke each stage of $A$ once, and $Q$ is the number of decryption queries made by $A$

### 7.3 Concrete instances

We shall show how to construct hash proof systems for certain symmetric subgroup membership problems. For completeness, we include proofs for this special case, even though the general case was done in [10].

Let $\mathcal{S S M}_{(G, K, H)}$ be a symmetric subgroup membership problem such that $G$ is cyclic, and suppose that a generator $g$ is available for $K$.

We shall describe a hash proof system $P$ and an extended hash proof system $\hat{P}$ for $\mathcal{S S M}_{(G, K, H)}$. Let $W=\mathbb{Z}_{|K|}$ and $\rho([w])=g^{w}$. Let $L=G$ and let $L^{\prime}=H$. Since $G$ is cyclic, the homomorphism group $\operatorname{Hom}(G, G)$ is isomorphic to $\mathbb{Z}_{|G|}$, and we let $M=\operatorname{Hom}(G, G)$. For any $f \in M$, a useful description of the function $\rho^{*}(f)$ is the group element $f(g)$, since for any $[w] \in W, f\left(g^{w}\right)=f(g)^{w}$. Then $P(G, K, H)=\left(G, \mathbb{Z}_{|H|}, \rho, \operatorname{Hom}(G, G)\right)$.

Proposition 7.2. The hash proof system $P$ described above is $\epsilon$-smooth, for $\epsilon \leq \frac{|H|-\phi(|H|)}{|H|}+\delta$, where the sampling algorithm for $M$ is $\delta$-close to uniform.
Proof. Apply Lemma 4.12.
The extended hash proof system $\hat{P}$ is slightly more complicated. The basic idea is to use the same construction as for $P$, but repeat it several times in parallel, and then combine the output. Let $\ell$ be the smallest prime factor in $|H|$. We shall suppose that for some sufficiently large $l$, a 1-1 function $h: G \times G \rightarrow\{0, \ldots, \ell-1\}^{l}$ is available.

Suppose we have a tuple $\vec{f}=\left(f_{0}, f_{1}, \ldots, f_{l}\right) \in \operatorname{Hom}(G, G)^{l+1}$. This tuple acts on $G^{l+1}$ in the obvious way. $h$ provides a 1-1 map $G \times G \rightarrow \mathbb{Z}_{|G|}^{l}$. We can extend this to a map $G \times G \rightarrow \mathbb{Z}_{|G|}^{l+1}$, given by $(x, e) \mapsto\left(1, \gamma_{1}, \ldots, \gamma_{l}\right)$. Since $\mathbb{Z}_{|G|} \simeq \operatorname{Hom}(G, G)$, we can let $h(x, e)$ act on $G^{l+1}$.

Now we compose the map $G \rightarrow G^{l+1}$ given by $x \mapsto(x, x, \ldots, x)$ first with $\vec{f}$, then with the action of $h(x, e)$, and finally with the $l$-fold multiplication map $\left(z_{0}, z_{1}, \ldots, z_{l}\right) \mapsto z_{0} z_{1} \ldots z_{l}$. This will be our $\hat{f}$. In other words, we consider the set of maps of the form

$$
\hat{f}(x, e)=f_{0}(x) \prod_{i=1}^{l} f_{i}(x)^{\gamma_{i}}
$$

where $h(x, e)=\left(\gamma_{1}, \ldots, \gamma_{l}\right)$, and $f_{i} \in \operatorname{Hom}(G, G)$.
The witness set for $K \times G$ is $\mathbb{Z}_{|K|} \times G$, and the map $\hat{\rho}$ is given by $\hat{\rho}([w], e)=$ $\left(g^{w}, e\right)$, where $g$ is a generator for $K$. It is clear that

$$
\hat{\rho}^{*}(\hat{f})([w], e)=f_{0}(g)^{w} \prod_{i=1}^{l} f_{i}(g)^{w \gamma_{i}}
$$

where $h\left(g^{w}, e\right)=\left(\gamma_{1}, \ldots, \gamma_{l}\right)$. So a useful description of the function $\hat{\rho}^{*}(\hat{f})$ is the tuple $\left(s_{0}, s_{1}, \ldots, s_{l}\right)=\left(f_{0}(g), f_{1}(g), \ldots, f_{l}(g)\right)$.

Proposition 7.3. The extended hash proof system $\hat{P}$ described above is $1 / \ell$ -universal-2.

Proof. We work with the definition (7.1).
Let $(x, e, \hat{y}),\left(x_{0}, e_{0}, \hat{y}_{0}\right) \in G \times G \times G,\left(k_{0}, \ldots, k_{l}\right) \in\{0, \ldots, l-1\}, h(x, e)=$ $\left(\gamma_{1}, \ldots, \gamma_{l}\right)$, and $h\left(x_{0}, e_{0}\right)=\left(\gamma_{0,1}, \ldots, \gamma_{0, l}\right)$.

Let $\left(x^{\prime}, x^{\prime \prime}\right),\left(x_{0}^{\prime}, x_{0}^{\prime \prime}\right),\left(\hat{y}^{\prime}, \hat{y}^{\prime \prime}\right)$ and $\left(\hat{y}_{0}^{\prime}, \hat{y}_{0}^{\prime \prime}\right)$ be the splittings of $x, x_{0}, \hat{y}$ and $\hat{y}_{0}$, respectively. Let $k_{i}^{\prime}=k_{i} \bmod |K|$ and $k_{i}^{\prime \prime}=k_{i} \bmod |H|$, for $0 \leq i \leq l$.

First of all, the knowledge of $s_{0}, \ldots, s_{l}$ fixes the values of $\left(k_{0}^{\prime}, \ldots, k_{l}^{\prime}\right)$. We may assume that

$$
\begin{equation*}
\hat{y}_{0}^{\prime \prime}=\left(x_{0}^{\prime \prime}\right)^{k_{0}^{\prime \prime}+\sum_{i=1}^{l} k_{i}^{\prime \prime} \gamma_{0, i}} \tag{7.9}
\end{equation*}
$$

since otherwise, (7.1) is trivially true. We would like to count how many tuples $\left(k_{0}, \ldots, k_{l}\right)$ satisfy this equation, and the proportion of those that satisfy

$$
\begin{equation*}
\hat{y}^{\prime \prime}=\left(x^{\prime \prime}\right)^{k_{0}^{\prime \prime}+\sum_{i=1}^{l} k_{i}^{\prime \prime} \gamma_{i}}, \tag{7.10}
\end{equation*}
$$

or alternatively, for some integer $a$,

$$
\begin{equation*}
a \equiv k_{0}^{\prime \prime}+\sum_{i=1}^{l} k_{i}^{\prime \prime} \gamma_{i} \quad(\bmod m), \tag{7.11}
\end{equation*}
$$

where $m$ is the order of $x^{\prime \prime}$. This will give us an upper bound on $\epsilon$.
It is clear that $\left(k_{0}^{\prime}, \ldots, k_{l}^{\prime}\right)$ and $\left(k_{0}^{\prime \prime}, \ldots, k_{l}^{\prime \prime}\right)$ are distributed uniformly and independently when $\left(k_{0}, \ldots, k_{l}\right)$ is sampled uniformly. Therefore, the only information about $\left(k_{0}^{\prime \prime}, \ldots, k_{l}^{\prime \prime}\right)$ available is in (7.9).

First of all, (7.9) just contains information about ( $k_{0}^{\prime \prime}, \ldots, k_{l}^{\prime \prime}$ ) modulo the order of $x_{0}^{\prime \prime}$. Let $m_{0}$ be the order of $x_{0}^{\prime \prime}$, and let $m^{\prime}=\operatorname{gcd}\left(m, m_{0}\right)$ and $m^{\prime \prime}=$ $m_{0} / m^{\prime}$.

Suppose $m^{\prime \prime}>1$. (7.11) also holds modulo $m^{\prime \prime}$, and it is quite clear that for any tuple $\left(k_{1}^{\prime \prime}, \ldots, k_{l}^{\prime \prime}\right)$, there is exactly one value of $k_{0}^{\prime \prime}$ that satisfies the congruence, so at most $1 / m$ of all tuples work. Since $\ell$ was the smallest prime dividing $|H|, 1 / m \leq 1 / \ell$, and the result follows.

Suppose $m^{\prime \prime}=1$, and let $\ell^{\prime}$ be the smallest prime dividing $m$. For some integer $a_{0}$, we get from (7.9) that

$$
\begin{equation*}
a_{0} \equiv k_{0}^{\prime \prime}+\sum_{i=1}^{l} k_{i}^{\prime \prime} \gamma_{0, i} \quad\left(\bmod \ell^{\prime}\right) \tag{7.12}
\end{equation*}
$$

The tuple ( $k_{0}^{\prime \prime}, \ldots, k_{l}^{\prime \prime}$ ) modulo $\ell^{\prime}$ is uniformly distributed in the affine space $\mathbb{F}_{\ell^{\prime}}^{l+1}$. (7.12) defines a hyperplane $W_{0}$ in $\mathbb{F}_{\ell^{\prime}}^{l+1}$, and (7.11) defines a hyperplane $W$.

Now each $\gamma_{i}$ and $\gamma_{0, i}$ is in $\{0, \ldots, \ell-1\}$, and since $h$ is $1-1$, for at least one $i, \gamma_{i} \not \equiv \gamma_{0, i}\left(\bmod \ell^{\prime}\right)$. This means that the two vectors $\left(1, \gamma_{1}, \ldots, \gamma_{l}\right)$ and $\left(1, \gamma_{0,1}, \ldots, \gamma_{0, l}\right)$ are linearly independent modulo $\ell^{\prime}$. Therefore, the two affine spaces $W$ and $W_{0}$ do not coincide and have non-empty intersection, and their intersection has codimension 1. Therefore, at most $1 / \ell^{\prime}$ of all tuples $\left(k_{0}^{\prime \prime}, \ldots, k_{l}^{\prime \prime}\right)$ satisfy both (7.9) and (7.10). Since $1 / \ell^{\prime}<1 / \ell$, the result follows.

Instantiating the cryptosystem $\Pi_{3}$ with these hash proof systems and the corresponding symmetric subgroup membership problem, we get the cryptosystem described in Figure 7.2. It is quite clear that if the underlying symmetric subgroup membership problem is hard, then by Theorem 7.1, this variant of $\Pi_{3}$ is secure, because the subgroup membership problem $\mathcal{S} \mathcal{M}_{\left(L, L^{\prime}\right)}$ that appears in Section 7.2 is simply $\mathcal{S M}_{(G, H)}$.

It is also possible to replace the 1-1 function $h$ with a collision resistant hash function, and take $l=1$. In this case, $\hat{P}$ is no longer $\epsilon$-universal-2, so we need to adapt the proof to show that $\hat{P}$ satisfies a "computational" variant of $\epsilon$-universal2. We do not do this.

Concretely, this cryptosystem can be instantiated with the symmetric subgroup membership problems $\mathcal{S O U}_{n}$ (Section 4.5.6) and $\mathcal{G O} \mathcal{U}_{n}$ (Section 4.5.7). Both yield quite practical and efficient cryptosystems, especially if a collision resistant hash function is used.

It is worth noting that if the extended hash proof system $\hat{P}$ is removed, we get the cryptosystem $\Pi_{2}$ from Section 5.3.

We note that for $\mathcal{S O} \mathcal{U}_{n}$, the group order $|G|=n$ is known. Sampling uniformly from $W=\{0, \ldots,|K|-1\}$ can be done using Lemma 2.3.

For $\mathcal{G O} \mathcal{U}_{n}$, we know that the uniform distribution on $\{0, \ldots,\lfloor n / 2\rfloor\}$ reduced modulo $|G|=\phi(n) / 2$ is almost uniform. Sampling uniformly from $W$ can be done using Lemma 2.3. We note that 2 divides the group order, and because of Proposition 7.2 and Proposition $7.3, K$ must be chosen such that 2 divides $|K|$.

If we consider performance, and compare $\Pi_{3} / \mathcal{S O} \mathcal{U}_{n}$ with ElGamal over $\mathbb{F}_{p}^{*}$, where $p$ is a safe prime, $\Pi_{3}$ requires slightly more than 4 exponentiations in $K$ to encrypt, while ElGamal requires 2 exponentiations in $\mathbb{F}_{p}^{*}$.

To decrypt, $\Pi_{3} / \mathcal{S O} \mathcal{U}_{n}$ requires less than three exponentiations in $K$ (under the assumption that the factorisation of $n$ is known to the secret key holder), while ElGamal requires 1 exponentiation in $\mathbb{F}_{p}^{*}$. If the factorisation of $n$ is not known, then decryption requires two exponentiations in $G$.

This means that $\Pi_{3} / \mathcal{S O} \mathcal{U}_{n}$ requires roughly the same amount of work as ElGamal to encrypt, and either 1.5 or 2 times as much work to decrypt. And it gives us chosen ciphertext security, not just semantic security as ElGamal does.

When instantiating with $\mathcal{G O}_{n}$, it is possible to make $K$ much smaller than $H$. This leads to an even faster system for encryption, but unless the factorisation of $n$ is known to the secret key holder, decryptions will cost between three and four exponentiations in $G$. If the factorisation is known, however, Chinese remainder theorem tricks are also available to speed up computations, giving a much faster system, even for decryption.

Compared to the instantiations of the Cramer-Shoup construction given in [10], our two instantiations are significantly faster, except for the elliptic curve variants of Cramer-Shoup. Asymptotically, the elliptic curve variants are significantly faster than our two variants.

## Key generation.

Input: $\mathcal{S S}_{\Gamma}, \tau$.

1. $(G, K, H) \leftarrow \Gamma_{\tau}$. A generator $g$ for $K$ is assumed given.
2. $\left(k, k_{0}, k_{1}, \ldots, k_{l}\right) \approx\{0, \ldots,|G|-1\}^{l+2}$.
3. $\left(s, s_{0}, s_{1}, \ldots, s_{l}\right) \leftarrow\left(g^{k}, g^{k_{0}}, g^{k_{1}}, \ldots, g^{k_{l}}\right)$.
4. $p k \leftarrow\left(G, g, s, s_{0}, s_{1}, \ldots, s_{l}, h\right)$.
5. $s k \leftarrow\left(G, k, k_{0}, k_{1}, \ldots, k_{l}, h\right)$.

Output: $(p k, s k)$.

## Encryption.

Input: $p k, m \in G$.

1. $w \approx\{0, \ldots,|K|-1\}$.
2. $x \leftarrow g^{w}$.
3. $y \leftarrow s^{w}$.
4. $e \leftarrow y m$.
5. $\left(\gamma_{1}, \ldots, \gamma_{l}\right) \leftarrow h(x, e)$.
6. $\hat{y} \leftarrow s_{0}^{w} \prod_{i=1}^{l} s_{i}^{w \gamma_{i}}$.

Output: $(x, e, \hat{y}) \in G \times G \times G$.

Figure 7.2: An instantiation of the cryptosystem $\Pi_{3}$.

When a key encapsulation method is all that is required, the Cramer-Shoup key encapsulation method [11] using a subgroup of a finite field will be significantly faster than our two constructions.

## Chapter 8

## Concluding remarks

Our presentation of subgroup membership problems, along with the related splitting and subgroup discrete logarithm problems, gives a uniform description and security analysis of a large class of public key cryptosystems, showing that they are all based on very similar principles.

Further, we have designed two key encapsulation mechanism based on a trapdoor splitting problem, and analysed their security in terms of the underlying splitting problem. To get security against chosen ciphertext attacks, we need the random oracle model and a strong assumption on the subgroup membership problem. Unfortunately, the concrete instances of this cryptosystem are in practice of moderate interest.

The most interesting part of this thesis is the study of symmetric subgroup membership problems. Theorems 4.11, 5.7 and 6.1 show how such an indistinguishability structure gives us very useful results.

Following the construction of Cramer and Shoup [10], we have designed a cryptosystem secure in the standard model against chosen ciphertext attacks. If a collision resistant hash function is used, the resulting cryptosystem is quite efficient when instantiated with a symmetric subgroup membership problem, and can compete with other instances of the Cramer-Shoup framework.

One problem is that, except for the Decision Diffie-Hellman problems, all of the subgroup membership problems we have presented depend crucially on factoring. While factoring is in general expected to be a hard problem, this dependency leads to relatively large groups when compared to cryptosystems based on elliptic curves. Finding subgroup membership problems, and especially subgroup discrete logarithm problems, that do not depend on factoring is a very interesting open problem.

## Bibliography

[1] M. Abdalla, M. Bellare, and P. Rogaway. DHIES: An encryption scheme based on the Diffie-Hellman problem, 2001. http://www.cs.ucsd.edu/ users/mihir/papers/dhies.html.
[2] M. Bellare, A. Boldyreva, and A. Palacio. An un-instantiable random-oraclemodel scheme for a hybrid-encryption problem. Cryptology ePrint Archive, Report 2003/077, 2003. http://eprint.iacr.org/.
[3] M. Bellare, A. Desai, D. Pointcheval, and P. Rogaway. Relations among notions of security for public-key encryption schemes. In H. Krawczyk, editor, Proceedings of CRYPTO '98, volume 1462 of $L N C S$, pages 26-45. SpringerVerlag, 1998.
[4] M. Bellare and A. Palacio. The knowledge-of-exponent assumptions and 3-round zero-knowledge protocols. Cryptology ePrint Archive, Report 2004/008, 2004. http://eprint.iacr.org/.
[5] M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Proceedings of the First ACM Conference on Computer and Communications Security, pages 62-73, 1993.
[6] J. Benaloh. Dense probabilistic encryption. In Proceedings of the Workshop on Selected Areas of Cryptography, pages 129-128, 1994.
[7] D. Boneh. The Decision Diffie-Hellman problem. In Proceedings of the Third Algorithmic Number Theory Symposium, volume 1423 of LNCS, pages 48-63. Springer-Verlag, 1998.
[8] D. Catalano, R. Gennaro, and N. Howgrave-Graham. The bit security of Paillier's encryption scheme and its applications. In B. Pfitzmann, editor, Proceedings of EUROCRYPT 2001, volume 2045 of LNCS, pages 229-243. Springer-Verlag, 2001.
[9] R. Cramer and V. Shoup. A practical public key cryptosystem secure against adaptive chosen cipher text attacks. In H. Krawczyk, editor, Proceedings of CRYPTO '98, volume 1462 of LNCS, pages 13-25. Springer-Verlag, 1998.
[10] R. Cramer and V. Shoup. Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In L. R. Knudsen, editor, Proceedings of EUROCRYPT 2002, volume 2332 of LNCS, pages 45-64. Springer-Verlag, 2002.
[11] R. Cramer and V. Shoup. Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. SIAM Journal of Computing, 33(1):167-226, 2003.
[12] I. Damgård. Towards practical public key systems secure against chosen ciphertext attacks. In J. Feigenbaum, editor, Proceedings of CRYPTO '91, volume 576 of $L N C S$, pages 445-456. Springer-Verlag, 1992.
[13] I. Damgård, J. Groth, and G. Salomonsen. The theory and implementation of an electronic voting system. In D. Gritzalis, editor, Secure Electronic Voting. Kluwer Academic Publishers, 2002.
[14] I. Damgård and M. Jurik. A generalisation, a simplification and some applications of Paillier's probabilistic public-key system. In K. Kim, editor, Proceedings of Public Key Cryptography 2001, volume 1992 of LNCS, pages 119-136. Springer-Verlag, 2001.
[15] Y. Desmedt and A. M. Odlyzko. A chosen text attack on the RSA cryptosystem and some discrete logarithm schemes. In H. C. Williams, editor, Proceedings of CRYPTO '85, volume 218 of LNCS, pages 516-521. SpringerVerlag, 1986.
[16] W. Diffie and M. E. Hellman. New directions in cryptography. IEEE Transactions on Information Theory, 22:644-654, 1976.
[17] D. Dolev, C. Dwork, and M. Naor. Non-malleable cryptography. In Proceedings of the 23rd Annual ACM Symposium on Theory of Computing, pages 542-552, 1991.
[18] T. ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. IEEE Transactions on Information Theory, 31:469-472, 1985.
[19] FIPS 180-1. Secure hash standard. Federal Information Processing Standards Publication. National Technical Information Service, Springfield, Virginia, April 1995.
[20] S. D. Galbraith. Elliptic curve Paillier schemes. Journal of Cryptology, 15(2):129-138, 2002.
[21] O. Goldreich. Foundations of Cryptography, Basic tools. Cambridge University Press, 2001.
[22] S. Goldwasser and S. Micali. Probabilistic encryption. Journal of Computer and System Sciences, 28:270-299, April 1984.
[23] S. Hada and T. Tanaka. On the existence of 3-Round zero-knowledge protocols. In H. Krawczyk, editor, Proceedings of CRYPTO '98, volume 1462 of $L N C S$, pages 408-423. Springer-Verlag, 1998.
[24] A. Joux. A one round protocol for tripartite Diffie-Hellman. In Proceedings of ANTS IV, volume 1838 of $L N C S$, pages 385-394. Springer-Verlag, 2000.
[25] H. W. Lenstra, Jr. Factoring integers with elliptic curves. Annals of Mathematics, 126:649-673, 1987.
[26] W. Mao. Fast Monte-Carlo primality evidence shown in the dark. Technical Report HPL-1999-30R1, HP Laboratories, October 1999.
[27] U. Maurer. Towards the equivalence of breaking the Diffie-Hellman protocol and computing discrete logarithms. In Y. Desmedt, editor, Proceedings of CRYPTO '94, volume 839 of $L N C S$, pages 271-281. Springer-Verlag, 1994.
[28] A. J. Menezes, P. C. V. Oorschot, and S. A. Vanstone. Handbook of applied cryptography. The CRC Press series on discrete mathematics and its applications. CRC Press, 1997.
[29] D. Naccache and J. Stern. A new public key cryptosystem based on higher residues. In K. Nyberg, editor, Proceedings of EUROCRYPT '98, volume 1403 of LNCS, pages 308-318. Springer-Verlag, 1998.
[30] M. Naor and M. Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In Proceedings of the 22nd Annual Symposium on Theory of Computing, 1990.
[31] J. M. G. Nieto, C. Boyd, and E. Dawson. A public key cryptosystem based on the subgroup membership problem. In S. Quing, T. Okamoto, and J. Zhou, editors, Proceedings of ICICS 2001, volume 2229 of LNCS, pages 352-363. Springer-Verlag, 2001.
[32] T. Okamoto and D. Pointcheval. The gap-problems: A new class of problems for the security of cryptographic schemes. In K. Kim, editor, Proceedings of Public Key Cryptography 2001, volume 1992 of LNCS, pages 104-118. Springer-Verlag, 2001.
[33] T. Okamoto and D. Pointcheval. REACT: Rapid enchanced-security asymmetric cryptosystem transform. In D. Naccache, editor, Progress in cryptology - CT-RSA 2001, volume 2020 of $L N C S$, pages 159-175. Springer-Verlag, 2001.
[34] T. Okamoto and S. Uchiyama. A new public-key cryptosystem as secure as factoring. In K. Nyberg, editor, Proceedings of EUROCRYPT '98, volume 1403 of LNCS, pages 308-318. Springer-Verlag, 1998.
[35] P. Paillier. Public-key cryptosystems based on composite degree residue classes. In J. Stern, editor, Proceedings of EUROCRYPT '99, volume 1592 of $L N C S$, pages 223-238. Springer-Verlag, 1999.
[36] C. Rackoff and D. Simon. Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In J. Feigenbaum, editor, Proceedings of CRYPTO '91, volume 576 of LNCS, pages 433-444. Springer-Verlag, 1992.
[37] R. L. Rivest, A. Shamir, and L. M. Adleman. A method for obtaining digital signatures and public-key cryptosystems. Communications of the ACM, 21:120-126, 1978.
[38] C. E. Shannon. Communication theory of secrecy systems. Bell System Technical Journal, 28(4):656-715, 1949. http://www.cs.ucla.edu/ ~jkong/research/security/shannon.html.
[39] V. Shoup. Lower bounds for discrete logarithms and related problems. In W. Fumy, editor, Proceedings of EUROCRYPT '97, volume 1233 of LNCS, pages 256-266. Springer-Verlag, 1997. Revised version available from http: //www.shoup. net.
[40] J. H. Silverman. The Arithmetic of Elliptic Curves, volume 105 of Graduate Texts in Mathematics. Springer-Verlag, 1986.
[41] A. Yamamura and T. Saito. Private information retrieval based on the subgroup membership problem. In V. Varadharajan and Y. Mu, editors, Proceedings of ACISP 2001, volume 2119 of LNCS, pages 206-220. SpringerVerlag, 2001.

## Index

$[a]_{n}, 3$
$\stackrel{\approx}{\leftarrow}, 29,32,38,54,56,63,68,71$, 76-78, 84
$\tau$, see security parameter
advantage, $7,8,11,18,22,30$
systematic, 35
$\operatorname{Adv}_{A}^{\mathcal{P}}$, see advantage adversary, 16
algorithm, 5
coin tosses, 6
cost, 6
efficient, 6, 7
generic, 24
polynomial-time, 7
probabilistic, 5
random tape, 5
rewinding, 6, 61
attack
active, 21
chosen ciphertext, $17,22,26$, 78,80
chosen plaintext, 16,62
non-adaptive chosen ciphertext, 17, 59
passive, 21, 71

Baby-step Giant-step, 44, 45
base, 10, 33
Bernoulli trial, 34
$k$-bit, 3

CCA1/2, see attack, (non-adaptive) chosen ciphertext
$\mathcal{C C} \mathcal{R}_{n}$, see Composite Residuosity problem
$\mathcal{C D}_{G}$, see Diffie-Hellman problem, Computational
challenge ciphertext, $16,26,59,70$, 73
Chinese remainder theorem, 5, 44
$\epsilon$-close, 4
complex multiplication, 49
Composite Residuosity problem, 46, $51,58,65,73$
cryptosystem
Cramer-Shoup, 83
ElGamal, 57, 62, 69, 83
Goldwasser-Micali, 57
homomorphic, 16, 20, 53, 62
hybrid, 21-23
key encapsulation method, 21-23, 67, 71, 85
Naccache-Stern, 58
Okamoto-Uchiyama, 58
Paillier, 58, 59
private key, see cryptosystem, symmetric
public key, 15, 53
symmetric, 21, 69
$\mathcal{D}$, see decryption algorithm
$\mathcal{D C} \mathcal{R}_{n}$, see Composite Residuosity problem
$\mathcal{D D}_{G}$, see Diffie-Hellman problem, Decision
decryption algorithm, $15,22,23$, $54,56,63,68,71,77,84$
symmetric, 21, 69
decryption problem, 17
DH-KEM, see Diffie-Hellman key encapsulation method

Diffie-Hellman key encapsulation method, 62, 73-74
Diffie-Hellman problem
Computational, 10, 45, 57, 70, 73
Decision, 11, 38, 45, 59, 70
gap, 73
discrete logarithm problem, 10, 33
distinguisher, 29
distribution, 3
almost equal, 4
multinomial, 13, 37
$\mathcal{D} \mathcal{L}_{G}$, see discrete logarithm problem
$\mathcal{E}$, see encryption algorithm
easy to compute, 6
elliptic curve, 5, 48, 49, 60, 70, 83
over rings, 5,44
encryption algorithm, 15, 22, 23, $54,56,63,68,71,77,84$
symmetric, 21, 69
encryptions indistinguishable from random, 19, 57
Euler phi-function, 3
expected value, 4
$\mathcal{F A C T}_{n}$, see factoring
factoring, $9,43,46,48,60,73$
elliptic curve method, 49
number field sieve, 49
find-stage, 16, 17
forger, 25
weak, 26, 61
Frobenius endomorphism, 5
generic model, 24, 59, 62
$\mathcal{G O U}, 49,50,51,58,60,73,83$
guess-stage, 16,17
hard, 9
hash function, 23, 70
collision resistant, 23, 83
hash proof system, 76, 80, 81
$\epsilon$-smooth, 78
$\epsilon$-universal-2, 78
extended, 76, 81, 82
Higher Residue problem, 43
higher residue problem, 58
$\operatorname{Hom}(G, G), 81$
$\mathcal{H} \mathcal{R}_{n}$, see Higher Residue problem
IND, see semantic security
indistinguishable encryptions, 18, 62, 78

Jacobi symbol, 42, 44, 50, 51
$J_{n}$, 42, 50
$\mathcal{K}$, see key generation algorithm
KEM, see cryptosystem, key encapsulation method
key generation algorithm, 15,22 , $54,56,63,68,71,77,84$
knowledge extractor, 26
weak, 26, 62
knowledge-of-exponents assumption, 59, 61
$\Lambda$, see cryptosystem, key
encapsulation method
Legendre symbol, 42, 57
lsd, 35
$\operatorname{Map}\left(S, S^{\prime}\right), 75$
negligible, 3
non-malleable, 20
one-way, 17, 55
oracle, 6
answer-tape, 6
decryption, 17, 79
encryption, 26
query-tape, 6
partial discrete logarithm problem, see subgroup discrete logarithm problem
perfect secrecy, 17
$\phi$, see Euler phi-function
$\Pi$, see cryptosystem
$p k_{C}, 15,22$
$p k_{K}, 22$
$p k_{M}, 15$
plaintext aware, 26
weakly, 26, 61
Pohlig-Hellman, 10, 35, 43
Pollard- $\rho$, 44, 50
polynomial security, see semantic security
private key, 15
probability space, 3
problem, 7, 29, 32, 33
equivalent, 8
hard, 8,38
parameterised, 8
reduction, 8
subproblem, 7, 29, 32, 33, 38
trapdoor, $9,32,33,65,67,71$
underlying, 9, 15, 32, 33
projection, 32
projection problem, see splitting problem
projective hash family, 75
public key, 15, 55, 57, 58
$Q_{n}, 42,43,49,50$
quadratic residue problem, 42,50 , 57
random function, 23
random oracle, 24
random oracle model, 24, 27
random variable, 4, 37
reduction, $7,34,55,57$
cost, 7
$\mathcal{R S} \mathcal{A}_{n}$, see RSA problem
RSA problem, 9, 46
safe prime, 11
$\mathcal{S D}$, see decryption algorithm, symmetric
$\mathcal{S D L}$, see subgroup discrete logarithm problem
$\mathcal{S E}$, see encryption algorithm, symmetric
security notion, 20
security parameter, 15,22
self-reducible, 7, 10, 11, 30, 32-35, 37
semantic security, 18, 21, 22, 57, 62, 67, 71
SHA- $i, 23$
shift cipher, 69
$\Sigma, 21$
significant, 3
$s k_{C}, 21$
$s k_{K}, 21$
$s k_{M}, 21$
$\mathcal{S M}$, see subgroup membership problem
smooth, 3
$\epsilon$-smooth, 76, 79, 81
Sophie-Germain-prime, 11
$\mathcal{S O U}_{n}, 49,60,69,83$
$\mathcal{S P}$, see splitting problem
splitting, 32
splitting problem, 32, 33, 53, 67
gap, 70, 71
hash, 74
oracle, 74
$\mathcal{S S M}$, see subgroup membership problem, symmetric
standard model, 23
statistical distance, 4
statistically indistinguishable, 4
subgroup discrete logarithm problem, 33-38, 55
subgroup membership problem, 29-31, 35, 38, 57, 62, 76, 80
associated, 76
symmetric, 38-42, 48, 49, 62, 67,81
success probability, 7, 8
trace, 8,26
trace-variant, 9, 26, 59, 61
Turing machine, 5
$\epsilon$-universal, 76
$\epsilon$-universal-2, $76,82,83$
witness set, 75


[^0]:    ${ }^{1}$ We follow Goldreich [21] and distinguish between information and knowledge. We have information about some value if we can compute the value given unbounded computational resources. We have knowledge about some value if we can compute it given bounded computational resources.

[^1]:    ${ }^{1}$ Historically, the non-adaptive attack was proposed first, and called chosen ciphertext attack, denoted by CCA. Then the adaptive attack was proposed, and people took to denoting the nonadaptive attack as CCA1 and the adaptive attack as CCA2.

[^2]:    ${ }^{1} \mathrm{~A}$ non-technical benefit is apparently avoiding a patent on the Paillier cryptosystem.

