



IST-2002-507932

ECRYPT

European Network of Excellence in Cryptology

Network of Excellence

Information Society Technologies

D.SPA.10

ECRYPT Yearly Report on Algorithms and Keysizes (2004)

Due date of deliverable: 31. January 2005

Actual submission date: 17. March 2005

Start date of project: 1 February 2004

Duration: 4 years

Lead contractor: Katholieke Universiteit Leuven (KUL)

Revision 1.1

Project co-funded by the European Commission within the 6th Framework Programme		
Dissemination Level		
PU	Public	X
PP	Restricted to other programme participants (including the Commission services)	
RE	Restricted to a group specified by the consortium (including the Commission services)	
CO	Confidential, only for members of the consortium (including the Commission services)	

ECRYPT Yearly Report on Algorithms and Keysizes (2004)

Editors

Christian Gehrman and Mats Näslund (ERICS)

Contributors

Steve Babbage (VOD), Dario Catalano (ENS), Louis Granboulan (ENS),
Arjen Lenstra (TUE), Christof Paar (RUB), Jan Pelzl (RUB),
Thomas Pornin (Cryptolog), Bart Preneel (KUL), Matt Robshaw (RHUL),
Andy Rupp (RUB), Nigel Smart (BRIS), Michael Ward (MasterCard)

17. March 2005

Revision 1.1

The work described in this report has in part been supported by the Commission of the European Communities through the IST program under contract IST-2002-507932. The information in this document is provided as is, and no warranty is given or implied that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability.

Contents

1	Introduction	1
1.1	Related Work	2
2	Algorithm Selection Criteria	4
3	Preliminaries	6
3.1	Notation	6
3.1.1	Primitive Information	6
3.1.2	Algorithm Information	7
4	Security Objectives and Attackers	8
4.1	The Security Process	8
4.2	Security Levels	9
4.3	Management Aspects	10
4.4	Attack Resources and a Note on Moore’s Law	10
4.5	Implementation Notes	11
I	General Key Size Recommendations	12
5	Determining Symmetric Key Size	13
5.1	PC/Software Based Attacks	14
5.2	Attacks Using ASIC Designs	14
5.3	Attacks Using FPGA Designs	14
5.3.1	Existing Area-Time Efficient FPGA Implementations	15
5.3.2	Exhaustive Key Search Based on FPGA Hardware	16
5.3.3	Cost Estimates	16

5.3.4	Research in the Field of FPGA-Related Attacks against DES	16
5.4	Conclusions	17
6	Determining Equivalent Asymmetric Key Size	18
6.1	Inter-system Equivalences	18
6.2	Survey of Existing Guidelines	19
6.2.1	Approach chosen by ECRYPT	21
6.3	Impact of Special Purpose Hardware	22
6.4	Quantum Computing	23
7	Recommended Key Sizes	24
7.1	Recommended Parameters: Non-confidentiality Objectives	25
7.1.1	Non-repudiation	25
7.1.2	Message Authentication	25
7.1.3	User Authentication	26
7.1.4	Hash Functions	27
7.1.5	Nonces	27
7.2	Security Levels	27
7.3	How to Deal with Very Long-term Security	27
7.4	A Final Note: Key Usage Principles	29
II	Symmetric Primitives	30
8	Block Ciphers	31
8.1	Overview	31
8.2	64-bit Block Ciphers	32
8.2.1	DES	32
8.2.2	3DES	32
8.2.3	Kasumi	33
8.3	128-bit Block Ciphers	33
8.3.1	AES	33
9	Stream Ciphers	34
9.1	Overview	34

9.1.1	A Note on Pseudo random Number Generation	35
9.2	RC4	35
10	Hash Functions	37
10.1	Overview	37
10.2	MD5	38
10.3	RIPEMD-128	38
10.4	RIPEMD-160	39
10.5	SHA-1	39
10.6	SHA224, SHA256	40
10.7	SHA384, SHA512	40
10.8	Whirlpool	41
11	Message Authentication Codes	42
11.1	Overview	42
11.2	HMAC	43
11.3	CBC-MAC-X9.19	43
11.4	CBC-MAC-EMAC	44
III	Asymmetric Primitives	45
12	Mathematical Background	46
12.1	Provable Security	46
12.2	Choice of Primitive	48
12.3	Non-cryptographic Attacks	48
13	Public-key Encryption	49
13.1	Overview	49
13.1.1	Security Notions	49
13.1.2	Consideration: Hybrid Encryption	50
13.2	RSA/Factoring Based	50
13.2.1	RSA PKCS#1 v1.5	50
13.2.2	RSA-OAEP	51
13.3	ElGamal/Discrete log based	52

14 Signatures	53
14.1 Overview	53
14.1.1 Security Notions	54
14.2 RSA/Factoring Based	54
14.2.1 RSA PKCS#1 v1.5	54
14.2.2 RSA-PSS	54
14.3 ElGamal/Discrete Log Based	55
14.3.1 DSA	55
14.3.2 ECDSA	56
15 Public Key Authentication and Identification	57
15.1 Overview	57
15.2 GQ	57
16 Key Agreement and Key Distribution	59
16.1 Overview	59
16.2 Factoring Based	60
16.2.1 RSA-KEM	60
16.3 DLOG Based	60
16.3.1 PSEC-KEM	60
16.3.2 ACE-KEM	61
16.3.3 Diffie-Hellman Variants	61
References	62
A Glossary	71

Abstract

This report contains the official delivery D.SPA.10 of the ECRYPT Network of Excellence, funded within the Information Societies Technology (IST) Programme of the European Commission's Sixth Framework Programme (FP6).

The report provides a list of recommended cryptographic algorithms (e.g. block ciphers, hash functions, signature schemes, etc) and recommended key sizes and other parameter settings (where applicable) to reach specified security objectives. The report reflects state-of-the-art in public knowledge at the time of writing. Due to possible advances in cryptanalysis, the report will be revised on a yearly basis.

The fact that a specific algorithm or variant thereof is not included in this report *should not* be taken as indication that particular algorithm is insecure. Reasons for exclusion could just as well be limited practical use (e.g. lack of standardization and/or implementation), maturity, etc.

Chapter 1

Introduction

To protect (information) assets in an IT system, cryptographic protocols, algorithms, and keys are used to reach certain security objectives deemed necessary for the protection of said assets. What characterizes today's approach is to rely on standardized, open "security frameworks" which are configured by "appropriate" algorithms, keys (and other parameters) to reach the security level as defined by the security objectives. Well-known examples of such frameworks are IKE, IPsec, TLS, S/MIME, etc. This is an important step forward from earlier approaches based on proprietary algorithms and protocols, kept (to the largest extent possible) unknown to the general public.

While it is recognized that the openness principle is the right way to go, it still does not make the problem of implementing security a trivial task since skill is still needed to determine which algorithms and keys are appropriate for the security objectives at hand. First of all, a better-safe-than-sorry approach, e.g. encrypting each message four times with different algorithms and huge key sizes, may not be the way to go, since it is likely to lead to bad performance, complex management, and in some cases bad "dependencies" between algorithms that could actually reduce security. As we will argue in more detail later, security is a process, rather than a state. In particular, the openness is sometimes a two-edged sword, leading to various attacks becoming known, and the average user may have difficulty in keeping up with cryptanalytic advances, *maintaining* a secure configuration of his/her system over time. Moreover, it is not always easy to understand the effects of *combinations* of different configuration options, e.g. what is the overall security level when protecting a k -bit AES key by an n -bit RSA key, using RSA padding option x ?

The purpose of this reports is to provide comprehensive, yet easy to use recommendations for the use of cryptographic algorithms, keys (key sizes), and other parameters in protocols such as those mentioned above. Specifically, the report contains the official delivery D.SPA.10 of the ECRYPT Network of Excellence, funded within the Information Societies Technology (IST) Programme of the European Commission's Sixth Framework Programme (FP6). While trying to keep the recommendations simple, we also at the same time provide extensive references for readers who are more technically oriented and would like to have more background information.

Since the field of cryptology is constantly advancing, the recommendations in this report will be updated at least on an annual basis. If a break-through result occurs, a special

report will be issued by ECRYPT, see e.g. [29]. The report reflects state-of-the-art in public knowledge at the time of writing.

The report is organized as follows. We conclude this chapter by comparing this report's role relative to other similar reports. In Chapter 2 we give rationale for algorithm selection criteria. Next, in Chapter 3 we introduce some notation and abbreviations used throughout the report. Then, in Chapter 4, we discuss security objectives on a high level, approaches to implement security relative to different type of attackers, and the importance of non-cryptographic issues, that are outside the scope of this report. The rest of the report is divided into three parts. Part I provides key size (and other parameter) recommendations. Part II provides recommendations on the use of symmetric algorithms, and Part III finally treats asymmetric algorithms.

1.1 Related Work

This is not in any way the first report containing recommendations for algorithms and key sizes. We will later survey previous work in more detail, but for the moment, we give a quick overview.

In [15], recommended key sizes for *symmetric* algorithms are given in relation to state-of-the-art in 1996. More recently, [70, 93] gives recommended key lengths for different attack scenarios and provides symmetric/asymmetric size-equivalence relations. A somewhat simpler version of the analysis in [70], geared towards the specific application of finding asymmetric sizes equivalent to (in practice fixed) symmetric sizes, can be found in [68]. These reports are quite generic and do not provide algorithm recommendations. In [121] a survey and comparison of various key-size recommendations can be found.

The US NIST recommends algorithms and key-sizes for US federal use through the publication of FIPS (Federal Information Processing Standard) documents—an algorithm/key-size appearing in a FIPS is thus considered “approved”.

The NESSIE consortium, [89], presents a portfolio of recommended algorithms, and in some cases, also key-size recommendations. Some of the algorithms in [89] have since been included in standards, some of which are covered here.

RSA Laboratories in [98] provide a cost-based analysis of equivalences between symmetric and asymmetric key sizes.

Finally, in [32], both algorithm and key-size recommendations can be found, but only for signature schemes.

In comparison to the above works, the current report aims at a wider scope, covering most types of algorithm primitives, in combination with recommendations for appropriate key sizes and the setting of other parameters. A second goal is to provide some easy to understand discussion on (security) properties of various primitives without using too technical language, yet striving to be correct. Indeed, experts in the area may find the text too informal in places (or even naive), so a third goal has been to provide numerous references to more technical treatments. As mentioned, the (at least) annual updates of this report aim to provide up-to-date information.

It is also important to distinguish between the questions *what is the smallest secure key size?*, and, *what is the smallest key size that is not completely insecure?* That is, rather than trying to find a magic “threshold” that separates security from insecurity, we have tried to be realistic and discuss more in terms of security levels, and elaborate considerations that could make a lower security level, forced onto us by environmental considerations in the form of bandwidth or processing power, acceptable. Also, the duration of the required protection is important, as is the value of the protected information asset.

Chapter 2

Algorithm Selection Criteria

The report provides a list of cryptographic algorithms in the following categories.

- Symmetric primitives:
 - block ciphers
 - stream ciphers
 - hash functions
 - message authentication codes.
- Asymmetric primitives:
 - public key encryption
 - signature schemes
 - public key authentication/identification schemes
 - key transport/key agreement.

Why are some algorithms included, yet other (sometimes well-known) algorithms excluded? The basis for selection has been security and wide-spread use. As a principle, only standardized, mature, wide-spread and *secure* algorithms are included. However, in some cases, commonly used, but security-wise non-optimal algorithms are also covered, pointing out caveats in using them. There are also a few cases where draft standards, anticipated to have substantial near-future impact have been included.

Therefore, the fact that a specific algorithm or variant thereof is *not* included in this report cannot be taken as indication that that particular algorithm is insecure. Reasons for exclusion may as mentioned also be limited practical use (e.g. lack of standardization and/or deployment), etc. For a survey of cryptographic standards, we refer to ECRYPT's report on this topic, [30]. Conversely, inclusion does not guarantee that a particular algorithm is secure, only that it is secure with respect to current state of the art. Current methods may fail, sometimes spectacularly, see e.g. [29].

This first yearly version of the report does not cover pseudo-random functions, pseudo-random generators, entity authentication using symmetric techniques or more compound

cryptographic protocols. Modes of operations for block ciphers are not treated either as the security is often well-understood, given an implementation using a secure block cipher, see e.g. [8]. In some cases “variants” of a basic algorithm, e.g. different RSA padding mechanisms, are covered due to their wide-spread use and/or expected advantages.

Chapter 3

Preliminaries

3.1 Notation

Throughout the report, the following notation is used.

$\log_b x$	logarithm of x to base b
$\log x$	$\log_2 x$
$\ln x$	$\log_e x$, $e = 2.7182\dots$
$ x $	the size (in bits) of x , i.e. $\lceil \log x \rceil$
\mathbb{Z}_m	the ring of integers modulo m
\mathbb{Z}_m^*	the multiplicative group $\subset \mathbb{Z}_m$
\mathbb{F}_s	the finite field of $s = p^n$ elements

Please refer to Appendix A for a glossary of abbreviations and acronyms.

We shall for each algorithm make references to various international standards making use of said algorithm. Though they are not standards in the formal meaning, the probably largest “consumer” of cryptographic standards are the IETF RFC specification suites. Since we shall often refer to these, we here wish to clarify use of some acronyms once and for all.

IPsec: IP security protocol, RFC 2401, 2402, 2406.

IKE: Internet Key Exchange, RFC 2409.

TLS: Transport Layer Security (TLS), RFC 2246.

S/MIME: Secure MIME, RFC 3396, 3850, 3851.

OpenPGP: RFC 2440, 3156.

These are available from <http://www.ietf.org/rfc.html>

3.1.1 Primitive Information

For each primitive type (block cipher, hash function, signature scheme, etc) we only give informal, intuitive definitions of what it means for the respective primitive to be “secure”. A

number of different security notions have been proposed and used in the literature, and it is beyond the scope of this report to investigate them more deeply, compare them, etc. When important for practical use, we shall highlight in slightly more detail what security notion we refer to. For the interested reader, we generally refer to the excellent survey given in each chapter of the NESSIE security report, [90]. Although our definitions are only informal, we shall not make use of any non-standard terminology, so [90] or any text book on cryptography, e.g. [78], should be able to provide more background, if needed.

3.1.2 Algorithm Information

Each of the included algorithms in Chapter 8 through Chapter 16, are represented in the form of an *algorithm record*, having the following shape and meaning:

Definition: Reference to stable algorithm specification.

Parameters: Parameters characteristic of the algorithm such as supported key and block sizes, as well as any data required to unambiguously specify algorithm operation, e.g. group used to carry out arithmetic etc.

Security: Statement about the security according to state-of-the-art. In particular, for symmetric algorithms, the effective key size, or, “as claimed”, meaning that no non-trivial attack is known. For asymmetric algorithms, where applicable, reference to proof of security and possible assumptions needed for the proof to hold.

Deployment: Reference to standards and/or products that use the algorithm.

Implementation: Pointers to “reference” implementation or test-vectors, if such are known to exist.

Public analysis: Reference to public analyses carried out, e.g. research papers and/or efforts such as NESSIE, Cryptrec, etc, where applicable.

Known weakness: Reference and short explanation of known weaknesses, should such be known.

Comments: Any additional information, e.g. caveats, pros/cons in using the algorithm, etc.

Note that the algorithm records are not necessarily exhaustive.

Chapter 4

Security Objectives and Attackers

We introduce security in IT systems to meet certain desired *security objectives*. Examples of well-known security objectives closely coupled to cryptography are *confidentiality*, *integrity*, and *non-repudiation*. The *security level* determines quantitatively to what extent these objectives need to be met, e.g. “how strong” confidentiality do we need, and is based on a threat and risk assessment of the IT system, which somewhat simplified asks questions of the following form:

1. How long into the future must security level persist? What is the “lifetime” and/or “value” (direct/indirect) of the protected asset?
2. Success of a brute-force key search is inevitable, but are there even better attacks?
3. What is the *attack model*? (Who is the attacker? What resources does he/she have?)
4. How will these resources develop during the lifetime of the protected data?
5. What cryptanalytic progress will occur during the lifetime of the protected data?

The approach taken in this report (and all other existing similar reports) is to take present state-of-the-art for the second point, make assumptions about the three last issues, and from that extrapolate key-sizes to match different security levels as defined by the first point. In the following we discuss in some more detail what this means.

Before doing so, a few words about defining the security level, i.e. (1) above. This is a business issue, figuring out what the overall financial impact of a security breach could be. Here, the more indirect value of “reputation” is typically a bigger consideration than direct monetary impact. This is then combined with the cost of various types (levels) of protections, and finally one decides what security level is desirable and till when.

4.1 The Security Process

It is important to realize that security is usually more of a process than a state. Security means to use and manage mechanisms to

Protect: make attacks expensive or unsuccessful.

Detect: make (attempted) attacks likely to be detected.

Respond: provide means to respond to successful attacks.

Recover: restore the system to a secure state after an attack.

Sometimes one also includes deterrence, which in a cryptographic sense is quite similar to protection (a large key size may deter attackers by the implied effort needed to break the system) but it is also covers non-technical issues such as legal frameworks to prosecute attackers. Such non-technical issues will not be discussed here.



Figure 4.1: The security process, [17].

One can identify two main approaches in implementing this security process. The *fail-safe* approach puts a lot of resources on deter/protect to make security failures unlikely, whereas what we (with slight abuse of english language) shall call the *safe-fail* approach spends effort more on detect/protect/recover, so that failures cannot be excluded, but can on the other hand be recovered from. As we will see, different security objectives are sometimes more suited for one of these two approaches.

Part of the goal of this report is to provide some guidenace on how to approach important parts of the security process in order to maintain a secure system. Still, many important aspects of the process are also out of scope, e.g. management processes surrounding cryptographic keys, such as revoking keys and replacing keys before they are “worn out” etc.

4.2 Security Levels

At the end of the day, what one usually cares about is how long it will take an attacker to “break” the security and what resources he needs in order to have a reasonable chance of succeeding. This cost (in time and/or money) of breaking the system must be higher than the value (in a similar measure) of the protected asset. For instance, if a 1Million dollar machine is needed to obtain secrets of value 1 dollar, one could hope the attacker is deterred by this (unless he/she simultaneously can get one million such secrets of course), and similarly, if information is only “sensitive” for a few hours, a break of the system requiring a week’s attack effort *might* be acceptable.

It is however important to note that the needed security level could depend on the security objective. Referring to the security process above, if confidentiality is compromised, there is

usually not much one can do to recover from the compromise. However, if we fear loss of non-repudiation to be close at hand, various combinations of recovery mechanisms may be available, e.g. re-signing with new algorithms and/or longer keys. Therefore, confidentiality is usually by nature a security objective that needs to be handled with the above defined fail-safe approach, whereas e.g. non-repudiation can sometimes also be handled by the safe-fail approach.

4.3 Management Aspects

When the security requirements are very high, it is usually the case that non-cryptographic issues become more important. For instance, to maintain an overall system security level of more than, say 128 bit symmetric keys, management and social aspects tend to weigh in more than mere key size. We do not claim that it is impossible or irrelevant to imagine very high security levels, but we want to stress that such non-cryptographic issues will then step in as the weakest link of the chain. This report therefore simply assumes that management issues can be handled to match the desired security. It is out of the scope of this report to convert such aspects into key sizes, attempting to answer a question like: “We use k -bit keys, managed by a process like this, what is the overall effective key size?”

4.4 Attack Resources and a Note on Moore’s Law

What attackers are realistic to take into account? In 1996, [15] used the following classification.

“Hacker”: using a \$0 budget and standard PC(s), or possibly a few \$100 spent on FPGA hardware.

Small organization: with a \$10k budget and FPGA(s).

Medium organization: \$300k budget and FPGA and/or ASIC.

Large organization: a \$10M budget and FPGA/ASIC.

Intelligence agency: \$300M budget and ASIC.

Such a classification is probably largely still valid today.

When security is to be maintained for longer periods than a few months, we must also take into consideration that the attacker may upgrade his/her resources according to developments in state-of-the-art. The sometimes debated, but commonly accepted way to handle this point is to assume Moore’s law. This law, today usually (mis)quoted as

$$\text{bits per square inch in year } t \sim 2^{(t-1962)/1.5},$$

has been a reasonably good approximation for the developments in “computing power per dollar” over the last few decades. For example, [15], mentions implementation of a DES

cracker on a certain family of FPGAs. The gate density of this family of FPGAs have increased from tens of thousands of gates (1995) to hundreds of thousands (2002).

Moore's formula is considered very inaccurate at describing CPU clock-speed, but for the probably more relevant measure of PC performance in terms of MIPS capacity, it seems to apply albeit with a slightly larger constant than "1.5", see [112] for a discussion.

Industry experts seem to agree that it is likely that Moore's law will continue to apply for at least a decade or more. Therefore we have chosen to adopt this assumption here (as is also done in e.g. [70] and previous key-size studies). However, completely new computational models and forms of hardware may also need to be considered, more on this later.

4.5 Implementation Notes

Cryptographic algorithms can be more or less sensitive to attacks that primarily exploit "physical" rather than cryptographic weaknesses. That is, attacks could be relying on implementation mistakes or exploitable environmental properties. For instance, some attacks are known based on e.g. a biased key-space, timing analysis, fault analysis, power consumption analysis, etc. Misuse and management errors can of course also lead to compromise. The recommendations in this report are given under the assumption that the algorithm is properly implemented, used, and managed, and furthermore run in an environment where it is not subject to the above mentioned side-channel attacks. Occasionally, however, we may point out specific caveats of this nature.

Part I

General Key Size Recommendations

Chapter 5

Determining Symmetric Key Size

For symmetric schemes, key size requirement is in principle quite straightforward. If the cryptosystem as such can be assumed to be secure for the lifetime of protected data, the only attack means is a brute force key search/guessing attack, whose time/success ratio only depends on the maximum amount of computing power (number of computers, special purpose hardware, etc), P , the foreseen attacker(s) have at their disposal. Thus, we select an n -bit key so that $\sim 2^n/P$ is somewhat larger than the life-time of the protected data. If the secret must prevail for more than, say a year, we should as mentioned also take Moore's law into account. Let us survey some earlier recommendations and some (publicly known) attack efforts, to benchmark how well extrapolation of old recommendations seem to hold.

Our basis will be the report [15], which in 1996 proposed 75-bit symmetric keys to protect "against the most serious threats", and 90-bits for a 20-year protection. The proposal is based on Table 5.1, which seems to be rationalized by defining minimum security as maintaining confidentiality a few days, up to a year. (The last column shows the assumed key-recovery time for the different attackers in the model used.) There are some indications that this is

Table 5.1: Minimum symmetric key-size in bits for various attackers (1996).

Attacker	Budget	Hardware	Minimum keysize	Recovery time
"Hacker"	0	PC(s)	45	222 days
	\$400	FPGA	50	213 days
Small organization	\$10k	FPGA	55	278 days
Medium org.	\$300k	FPGA/ASIC	60	256 days
Large org.	\$10M	FPGA/ASIC	70	68 days
Intelligence agency	\$300M	ASIC	75	73 days

still a reasonably accurate approach. We have some recent data-points of attack efforts, as well as some proposed hardware (FPGA and ASIC) designs to compare to.

5.1 PC/Software Based Attacks

In late 1999, a small team which for these purposes could be considered (skilled) “hackers” found a 48-bit DES key in about 3 weeks using a relatively small number of general purpose computers, distributing key-search as part of a screen-saver application, [3]. Applying Moore’s law to Table 5.1, it predicts that 48-bit keys in 1999 would resist hackers for about a year. Thus, the attack was about 17 times faster than expected. On the other hand it is doubtful if [15] included such an attack in the “hacker” category, and we return to this later.

In 2002, [28] completed the RC5 64-bit challenge after nearly 5 years and some 300,000 individual general-purpose computers participating. (The same source completed the 56-bit challenge in 250 days in 1997 using 3,000 machines.) A question is how to interpret such results, since they are done by “hackers” with relatively small resources, utilizing the connectivity offered by today’s (and tomorrow’s) Internet. This is an important issue, since if one indeed considers this a hacker attack, the hacker key sizes need a quite big increase relative to predictions by the table (at least 5-10 bits). One could on one hand argue that such resources will only be summoned on “challenge” types of attacks. On the other hand, various forms of Malware (e.g. worms¹) are more and more frequent, and the average user could, unbeknownst to him/her, be part of a real attack. Indeed, [107] reports that this is one of the most popular applications of worms. These case-studies leads us to believe that we need to introduce a new type of hacker using distributed attacks, perhaps stealing CPU time from victims at least on a medium scale, and we propose to add an additional 8 bits for this category. We shall not explicitly consider the very large scale challenge type efforts.

5.2 Attacks Using ASIC Designs

EFF’s DES-cracker (a.k.a. “Deep Crack”), [36], designed and built at a cost of about US\$200k, was reported to recover 56-bit DES keys in just over 22 hours in 1998. Table 5.1 predicts that in 1998, keys of roughly 61 bits would resist such attackers for 2 days. Extrapolating, Deep Crack would have retrieved such keys in about 1 month. Table 5.1 is thus slightly conservative, but close.

In [18] (1999), a \$280M machine is sketched that would occupy an area one fifth of the Pentagon and would retrieve 64-bit RC6 keys in a few minutes. It is noted that this figure matches quite well an extrapolation of [15] numbers.

5.3 Attacks Using FPGA Designs

In the following, we will estimate the costs for an exhaustive key search attack on the symmetric ciphers DES and AES based on the use of FPGA devices. As basis for our estimates, we use area-time optimized FPGA implementations of DES and AES presented in the next section. As main result of our assessment we conclude that a DES key search using FPGAs is

¹A piece of computer program that automatically spreads to hosts on a network, infecting them with “malware”.

very feasible at low costs and moderate engineering expertise. In the case of AES, however, a brute force attack seems completely out of reach for several decades to come.

The reason why we focus in more detail on FPGAs as opposed to ASICs (Application Specific Integrated Circuits) is that that FPGAs require (1) considerably less engineering expertise and (2) the initial costs for an FPGA development environment are very moderate (perhaps a few 1000s of dollars vs. 100,000s of dollars in the ASIC case, see, e.g. [31]). We would like to stress, though, that ASICs are more cost efficient than FPGAs in high volume applications. Thus, in the case of key search machines for AES, one could possibly achieve lower costs with dedicated ASICs. However, the expected cost/performance benefit should be not more than 1–2 orders of magnitude, so that even an ASIC-based key search attack seems completely out of reach.

5.3.1 Existing Area-Time Efficient FPGA Implementations

A comparison of different AES implementations on FPGAs can be found in [45]. Table 5.2 summarizes relevant information and the corresponding references.

Table 5.2: Comparison of some published FPGA AES implementations (128 bit key)

Design	Device	Slices	Block RAM	Throughput (Gbit/s)	Mbit/s / Slice
Gaj et al. [34]	XCV1000-6	12600	80	12.1	0.96
McLoone et al. [72]	XCV812E-8	2222	100	6.95	3.1
Standaert et al. [110]	XCV3200E-8	2784	100	11.77	4.2
Saggese et al. [99]	XVE2000-7	5810	100	20.3	3.4
Hodjat et al. [45]	XC2VP20-7	5177	84	21.54	4.0

Apparently, [110] and [45] yield the best performance of AES on an FPGA regarding the area-time trade off. A throughput of 11.77 Gbit/s and 21.54 Gbit/s could be reached, respectively. Another contribution shows efficient FPGA accelerators both for DES and AES ([64]). For a throughput of 744 MByte/s, DES requires 3,250 CLBs (Configurable Logic Blocks) on a single FPGA. For a throughput of 2,500 MByte/s, AES requires 5,350 CLBs and 80 Block RAMs.

Besides results from the research community, there are also commercial cores for symmetric ciphers available, e.g., from Helion. The “Helion AES Core” runs on a low-cost FPGA (Xilinx Spartan-3) and is available in four different gradings: a “Tiny Core” with a throughput of 18 Mbit/s, a “Standard Core” with 229 Mbit/s, a “Fast Core “ with approx. 1 Gbit/s, and a “Pipelined Core” with more than 10 Gbit/s ([44]).

For DES, Helion offers a core running on a Xilinx FPGA (Virtex-II) with a data rate greater than 640Mbit/s for single DES and more than 230Mbit/s for triple DES ([44]).

5.3.2 Exhaustive Key Search Based on FPGA Hardware

For an exhaustive key search, all possible keys are tested with the FPGA hardware implementations of DES or AES. Assuming the possession of a pair clear-cipher text, we only need to encrypt a single block with each key candidate. All computations are based on the results of [110, 64], where non-pipelined implementations are used (pipelining cannot be used since we only encrypt/ decrypt a single block with one specific key).

For AES, we can test approximately $9.2 \cdot 10^7$ keys per second. With the DES hardware, $1.2 \cdot 10^7$ keys per second can be checked. Hence, with a single encryption or decryption unit, an exhaustive key-search would take on average $3.1 \cdot 10^9$ s (≈ 98 years) and $1.9 \cdot 10^{30}$ s ($\approx 5.9 \cdot 10^{22}$ years) for DES and AES (128 bit key), respectively. These numbers are already a first indication of the feasibility of an exhaustive key search for DES vs. AES.

5.3.3 Cost Estimates

In order to achieve low over-all costs, we assume high volume prices for FPGAs. For instance, a typical low-cost FPGA such as the Xilinx Spartan-3 FPGA (90nm process) currently costs about \$6.50 ([122]) and contains approximately 8,000 configurable logic blocks and, hence can fit two AES or two DES units including overhead. Since the FPGAs listed in Section 5.3.1 differ from the Xilinx Spartan-3 platform, the following cost calculations are not exact but should give a valid estimate.

Assuming an exhaustive key search within one month, many FPGAs have to be assembled in a cluster, each performing key searches in a part of the key space. We assume an overhead of 100% for printed circuit boards, power supplies etc. Costs for development, electrical energy and possibly cooling are not included in this estimate. Under these conditions, DES can be cracked in one month with 75 FPGAs (containing two DES engines each), at a cost of \$13 (FPGA plus 100% overhead), yielding a total cost of approximately \$1,000.

Note that these numbers are based on an average key search, i.e., half the key space is searched. The cost-performance ratio stays constant in this estimation. That means, for instance, that a ten times faster key search (i.e., 3 days) would require a key search machine which is ten times more expensive (i.e., \$10,000). Again, we would like to stress that these are estimates rather than exact numbers.

For AES (128 bit), a machine performing an average key search in one month would cost $\$4.6 \cdot 10^{24}$. Even if we take Moore's Law into account (i.e., decrease of IC costs by a factor of two every 1.5 years at constant performance), such a machine would still cost $\$4.2 \cdot 10^{14}$ in the year 2055. It is, however, somewhat doubtful if Moore's Law will be valid for such a time span, as semiconductor-based computers will run in various physical limitations before that. However, this is also a cost related issue: even if technology gets stuck, manufacturing capacities may still increase and costs decrease, thus keeping Moore on track, possibly.

5.3.4 Research in the Field of FPGA-Related Attacks against DES

A paper by Hamer et al. ([41]) describes a DES-cracking hardware on a field-programmable system called the "Transmogripher 2a". A fully implemented system would be able to search

the entire key space in 1040 days at a rate of 800 million keys/second. A single unit consists of two Altera 10K100 FPGAs, each connected to up to 4MB memory. The whole system consists of 16 such units and was expected to cost approximately \$60,000.

Besides a simple exhaustive key search machine, [66] presents a hardware FPGA implementation of linear cryptanalysis of DES. The authors implemented a modified attack from Matsui's original attack. The resulting attack is less efficient than Matsui's attack, but fits in the hardware and breaks a DES key in 12-15 hours on one single FPGA.

A quite accurate estimate of an attack on DES by an exhaustive key search is given in [109]. Assuming the use of low-cost devices such as the 3S1000 Spartan-3 from Xilinx (\$ 12 p.p.) and recent results from efficient DES implementations, a device for cracking DES within 3 days would cost approximately \$ 12,000.

5.4 Conclusions

Motivated by the analysis above, our proposal is to (according to Moore's Law) add 5-6 bits to all the figures in Table 5.1, and to add the distributed attack type as above with some additional safety margin. ECRYPT's symmetric key-size table is found in Chapter 7.

Let us do a quick "sanity check" by comparing the above conclusion to the work in [70]. That study is based on a key size requirement defined in terms of Moore's Law and that 56-bit (DES) keys were "sufficiently infeasible" in terms of a 1982 \$50M machine retrieving 56-bit keys in 2 days. Now, for 2004 [70] recommends symmetric keys of about 73 bits. Our extrapolation (to add 5-6 bits to [15] values) here leads to a (2004) \$300M machine finding 81-bit keys in 73 days. It is tempting to turn the (\$50M, 73-bit, 2-day) machine into a \$300M machine that recovers (almost) 76-bit keys in the same time, or, (almost) 81-bit keys in 64 days. This thus compares well with our estimate.

For FPGAs, Table 5.1 predicts that in 2004, 56-bit keys will give about 200 days of protection against a \$400 FPGA design. Above, we estimated that \$12,000 would give us a device that retrieves such key-sizes in 3 days. Scaling, \$400 spent on FPGA would give us the 56-bit keys if we are willing to wait 30 times longer, i.e. 90 days.

Chapter 6

Determining Equivalent Asymmetric Key Size

For asymmetric cryptography, things are more complicated since

- unnecessarily large keys affect performance
- increasingly effective attacks (all *much* better than brute force) have been discovered over the last 30 years
- we have recently seen suggestions for special purpose cryptanalytic hardware that are more than being “naive” search-machines.

Symmetric and asymmetric schemes are often used in conjunction, e.g. an asymmetric key is used to guard a symmetric key. Thus, we must be aware of the somewhat worn-out, yet universally valid “security is as good as the weakest link” principle, and we need a way to find an asymmetric key-size that matches a pre-specified symmetric key-size. At the same time, we want to avoid an overly large public key. A number of comparative studies have been made by skilled people, and rather than doing it again, we would like to see what conclusions can be drawn from these sources.

6.1 Inter-system Equivalences

In the following, *RSA* refers generically to a public key scheme, assumed equivalent in hardness to the integer factoring problem. We similarly use *DLOG* and *EC* to refer generically to schemes based on discrete logarithms and (secure) elliptic curve groups, respectively, both over finite fields of prime order.

For EC, state-of-the-art suggests that EC in a subgroup of m -bit size is equivalent to $m/2$ bits symmetric key. In fact, considering that the elementary operations of these attacks are point additions on an elliptic curves, which should be compared to trying a symmetric key, it would be possible to reduce EC key size by roughly 8-12 bits for practical ranges of key-sizes. This is done in [70], but not in e.g. [84, 68]. We propose to follow the simple half-the-size principle as a general recommendation. This will also, as noted in [70], provide some

protection against unforeseen developments in cryptanalysis of elliptic curves, without serious impact on performance. Curves over binary fields are often recommended (see e.g. [120]) to use slightly larger keys to mitigate certain hardware attacks. Also, the generic attacks are in the binary case sometimes a bit more efficient. For instance, so-called Koblitz curves over binary fields of size 2^n can be attacked roughly a factor $\sqrt{2n}$ faster. Curves over binary fields generally also require other forms of special attention when selecting parameters.

According to state-of-the-art, the difficulty of solving DLOG in prime order fields of size 2^n is, up to constants, asymptotically equivalent to that of breaking n -bit RSA. In practice though, DLOG is noticeably more difficult. Moreover, DLOG is in most standardized algorithms performed in a smaller *subgroup*, and then the size of this subgroup that matters too, in which case the symmetric key equivalent is theoretically half of the bit-size of said subgroup, again according to the same generic attacks applicable also in the EC case. This would imply DLOG sub-groups of the same size as a corresponding EC group. Note though that performing exponentiations over a finite field is noticeably more expensive than on an elliptic curve of equivalent security. (The difference can be on the order 10-40 times, depending on security level.) Implementations which are concerned with performance could therefore, if required, reduce subgroup size by a few bits without lowering the security compared to the EC case (e.g. [70, 93] uses roughly 10-bits size reduction of the subgroups for practical key-sizes). Nevertheless, our general recommendation shall for simplicity be according to the same half-the-size principle, and to assign n -bit fields the same security as n -bit RSA, a slightly conservative recommendation for DLOG based schemes.

With this in mind, the main issue is to determine RSA vs. symmetric key equivalence. We next survey some existing publications in this area.

6.2 Survey of Existing Guidelines

Assuming that the 56-bit Data Encryption Standard (DES) was considered “secure” by any reasonable meaning until 1982, [70] suggest that 80-bit symmetric keys are secure until 2012 and are *computationally* equivalent to 1120-1464 bit RSA/DLOG keys (depending cost model; HW/general purpose computer), 140 bit DLOG subgroups and 149-165 bit EC groups (depending on likelihood of cryptanalytic progress). The analysis is quite methodical, and can be used to derive key-sizes relative to varying assumptions and scenarios.

Table 6.1: Lenstra-Verheul recommendations (assuming “average” cost model/cryptanalytic progress).

Equivalent symmetric key size	56	64	80	96
Elliptic curve key size	95	116	160	200
Modulus length (pq) /dlog field	380	580	1300	2500
Dlog subgroup	102	114	141	171

Taking *cost* (of memory) more into consideration, [106] argues that in the year 2000, 1024-bit RSA keys corresponded to 96-bit symmetric keys. As noted in [103], this leads to quite different extrapolations of the lifetime of RSA keys between [70] and [106], the second giving

1024-bit RSA keys some 30 years extra lifetime. On the other hand, worst-case scenarios in [70] suggests that 1024-bit RSA keys could already have been factored today, but there is no public record that any key even close in size has been factored.

In the US, NIST has issued an equivalence table, [84], stating that 80-bit symmetric keys correspond to 1024-bit RSA/DLOG keys and 160-bit DLOG subgroups and EC groups. For the 128-bit level, 3072 and 256-bit keys, respectively, are recommended. The 80/1024/160 bit level is considered sufficient until 2010, beyond that 112/2048/224 bits are promoted. RSA

Table 6.2: NIST recommendations.

Equivalent symmetric key size	80	112	128	192	256
Elliptic curve key size	160	224	256	384	512
Modulus length (pq)/dlog field	1024	2048	3072	7680	15360
Dlog subgroup	160	224	256	384	512

Lab’s and Certicom’s recommendations are in line with this. For all practical purposes, this is also in reasonable agreement with [70] but does not differentiate the cost of operations in EC and DLOG subgroups, compared to symmetric key operations.

Recently the NESSIE consortium, in [90], for “medium term” (5-10 years) security recommends the use of 1536-bit keys for RSA and DLOG based public key schemes, and 160-bit for elliptic curve discrete logarithms, suggesting a 1536/80 equivalence, which is in line with [70], with the same cost-remark for EC as above. This recommendation is based on an as-

Table 6.3: NESSIE recommendations.

Equivalent symmetric key size	56	64	80	112	128	160
Elliptic curve key size	112	128	160	224	256	320
Modulus length (pq)	512	768	1536	4096	6000	10000
Modulus length (p^2q)	570	800	1536	4096	6000	10000

sumed equivalence between 512-bit RSA keys and 56-bit keys, and an extrapolation of that. However, it should be noted that during the course of producing this report, we discovered that the NESSIE recommendations were based on a slightly inaccurate extrapolation formula, leading to somewhat too large RSA keys.

RSA Labs, in [98], performs a cost-based analysis, and arrives at the following key-size equivalences (Table 6.4). The time to break is computed assuming a machine that can break

Table 6.4: RSA Labs analysis.

Symmetric Key	EC	RSA	Time to Break	Machines	Memory
56	112	430	< 5 min	105	trivial
80	160	760	600 months	4300	4Gb
96	192	1020	$3 \cdot 10^6$ yrs	114	170Gb
128	256	1620	10^{16} yrs	.16	120Tb

a 56-bit DES key in 100 seconds, then scaling accordingly. The Machines column shows how

many NFS sieve machines can be purchased for \$10 million under the assumption that their memories cost \$.50/Mbyte¹.

The IETF recommendation RFC3766 [93] (which is largely based on the cost-model variation of [70]) suggests that an 80-bit symmetric key is equivalent to 1228-bit RSA/DLOG keys, and 148-bit DLOG subgroups. Specifically, Table 6.5 is given. IETF also issued rec-

Table 6.5: IETF RFC3766 recommendations.

Equivalent symmetric key size	80	100	150	200	250
Modulus length (pq) /dlog field	1228	1926	4575	8719	14596
Dlog subgroup	129	186	284	383	482

ommended lifetime for (short) RSA signatures, [119]. This is based on a 100-fold security

Table 6.6: Short RSA signature key lifetimes from [119].

Keysize	Lifetime
512	1 hour
576	10 hours
640	4 days
704	30 days
796	8 months
1024	1 year

margin, e.g. the 512-bit keys should required 5 days of attack effort.

Basically all of the above reports explicitly or implicitly discuss requirements for *confidentiality*. It is sometimes argued that requirements for *authenticity* can be lower since data can be re-authenticated at regular intervals with increasingly large keys and even new algorithms. For signatures, the ETSI report [32] “approves” key-sizes inline with those above: 1024-bit minimum factoring/finite field discrete log keys and 160-bit elliptic curve groups size, but the approval is to be re-evaluated after five years (i.e. the approval is valid for the period 2001-2005).

6.2.1 Approach chosen by ECRYPT

Let n_{512} be the symmetric (e.g. DES) equivalent key-size to a 512-bit RSA key. Similar to NESSIE and motivated by the above discussion, ECRYPT’s public key size recommendations for factoring/RSA based and DLOG based schemes are based on an estimate for n_{512} and an extrapolation of the attack complexity of the general number field sieve algorithm (GNFS), the presently fastest method for computing integer factorization and discrete logarithms.

Specifically, the run-time of GNFS to factor N is estimated to (asymptotically) be

$$L(N) = Ae^{(C+o(1))(\ln N)^{1/3}(\ln \ln N)^{2/3}},$$

¹The .16 entry at the 128-bit level signifies that only .16 of the needed 120Tb memory can be purchased for the available \$10M budget.

for a constant A , and $C = (64/9)^{1/3}$. We shall make the assumption that the $o(1)$ -term can, for the key size ranges at hand, be treated as zero. From this, we can calculate A based on $L(512) = n_{512}$.

This leaves us with the problem of determining the quantity n_{512} . Experience from available datapoints suggests that the “resistance” of RSA-512 is some 4 – 6 bits less than that of DES-56. We here choose the more conservative end of this interval, i.e. $n_{512} = 50$.

We now get the following expression for the effective key-size of n -bit RSA modulus N :

$$s(n) = \left(\frac{64}{9}\right)^{1/3} \log_2(e)(n \ln 2)^{1/3} (\ln(n \ln 2))^{2/3} - 14.$$

For elliptic curves and discrete log subgroups, as discussed above, we apply the “half-the-size” principle.

6.3 Impact of Special Purpose Hardware

For a (secure) symmetric primitive, the only means to speed up attack time is to use parallelism by running on several general-purpose computers, or, to build a special purpose key-search machine. To our knowledge, the largest such machine built, is the EFF’s “Deep Crack”, [36], mentioned above. Such hardware is a reality one has to consider, and even without further improvements, a 64-bit key would be recovered in less than a year.

Also for public key (number theoretic) schemes, special purpose cryptanalytic hardware has been proposed.

The TWIRL device and its (optical) TWINKLE predecessor, [100, 101], have been proposed to speed up the so called sieving-step of factoring algorithms. It seems to be acknowledged (at least by RSA Labs, [63], 2003) that TWIRL would seriously threaten 512-bit RSA keys (which we already know can be attacked by general purpose computers and a quite small effort, [3]) and that the cost of building a TWIRL device that factors 1024-bit RSA keys “reasonably fast” is comparable to the cost of a parallel key-search machine retrieving 80-bit symmetric keys in comparable time.

Bernstein, [12], proposes a massively parallel machine for speeding up the so-called matrix step of factoring algorithms and claims an impact of factoring based key sizes, requiring three times as large keys. Later analysis [69], however, suggests a more moderate 17% increase, and only under “optimistic” assumptions.

Finally, [91] (1999), describes a hardware design estimated to cost \$10M that could possibly break elliptic curves over a binary field of size 2^{155} in about one month. However, since the design would be less cost-effective for general finite fields, a way to avoid issues is to avoid binary fields. In [120] it is estimated that curves over binary fields should have 8-10 bit larger keys to mitigate the effect of such hardware.

Considering that most of the proposed machines above seem to effect only a limited range of key sizes, or have limited impact, for the purpose of this report, the only special purpose hardware we shall consider (besides hardware for finding symmetric keys) is the machine discussed in [91]. Thus, we propose concerned users to add about 10 bits in the binary field

case. However, it is important to keep up to date with advances in special purpose hardware; we by no means rule out future advances in this area.

6.4 Quantum Computing

Both of the fundamental intractability assumptions on integer factoring and discrete logarithms break down if a (large) quantum computer could be built as demonstrated by Shor, [104]. For instance, integers N can be factored in only $O(\log^3 N)$ “steps” on such a machine. We are, however, quite far from realizing such a device. In [113], an experimental result for factoring the “toy” number $N = 15$ is reported with a run-time of just under one second.

For symmetric cryptography, the effect would also be dramatic, though not devastating. By the generic search algorithm due to Grover, [38], key-sizes are in effect cut in half. Also this algorithm has been implemented on small toy examples, [22].

The recommendations in this report assumes (large) quantum computers do not become a reality in the near future.

Chapter 7

Recommended Key Sizes

With reference to the above discussion, ECRYPT recommends the following minimum key sizes to protect against different attackers. Note that these are *minimum* sizes, giving pro-

Table 7.1: Minimum symmetric key-size in bits for various attackers.

Attacker	Budget	Hardware	Min security
“Hacker”	0	PC	51
	< \$400	PC(s)/FPGA	56
	0	“Malware”	59
Small organization	\$10k	PC(s)/FPGA	62
Medium organization	\$300k	FPGA/ASIC	66
Large organization	\$10M	FPGA/ASIC	76
Intelligence agency	\$300M	ASIC	81

tection only for a few months. The sizes need to be scaled accordingly, depending on the security level, see Section 7.2.

Given any desired (symmetric key) security level, one may need to convert the symmetric key into an equivalent asymmetric key size. Based on the preceding discussion, we propose the following symmetric/asymmetric size-equivalences, see Table 7.2. We note that it agrees reasonably well with Table 1 of [68].

It may also be interesting to see what today’s commonly deployed RSA/DLOG key-sizes offer in terms of equivalent symmetric key-size. This can be found in Table 7.3.

Note that the DLOG and EC recommendations applies for prime order fields. For finite field DLOG schemes, the table may need to be revised taking into account more efficient attacks that are known to exist for DLOG over binary fields. However, we are not aware of any practical deployment of systems based on discrete logarithms in binary fields.

For EC, recommendations for binary fields need to take into account that the field size should be 2^p , where p is a prime number¹ of bit-size slightly larger than the group/key bit-size. If in addition, the special purpose HW of [91] is considered a threat, roughly an additional

¹Unless p is a prime, more efficient attacks are known, [118].

Table 7.2: Key-size Equivalence.

Security (bits)	RSA	DLOG		EC
		field size	subfield	
48	480	480	96	96
56	640	640	112	112
64	816	816	128	128
80	1248	1248	160	160
112	2432	2432	224	224
128	3248	3248	256	256
160	5312	5312	320	320
192	7936	7936	384	384
256	15424	15424	512	512

Table 7.3: Effective Key-size of Commonly used RSA/DLOG Keys.

RSA/DLOG Key	Security (bits)
512	50
768	62
1024	73
1536	89
2048	103

10 bits should be added to the EC key size in this case.

7.1 Recommended Parameters: Non-confidentiality Objectives

The above minimum requirements are intended to keep messages' confidentiality for a certain time. For other security objectives, things may be a little different.

7.1.1 Non-repudiation

The possibility to recover security is better here, since one can in principle re-authenticate (and possibly revoke) at regular intervals when one suspects an imminent compromise. Therefore, somewhat shorter keys *may* be acceptable.

7.1.2 Message Authentication

Basically, the same recovery mechanisms may be applicable here. Secondly, some applications are of such nature that the value of the data as "authentic" is essentially only a real-time issue, and efficiency/bandwidth may in addition here require a lowered security.

Since there is usually no or little gain (in efficiency) in using short (symmetric) keys, the above recommendations hold (in general, see discussion below) also for authentication key sizes. Assuming a large enough key is chosen (and secure protocols/functions are used), the main threat is therefore forging by “guessing”. (We note that there are MACs, e.g. Carter-Wegman based MACs, where one can guarantee that such attacks are the only ones possible.)

For MAC tags, it may be acceptable to use shorter values than the MAC key size. In [42] an absolute minimum tag size of 32 bits is motivated. In [6] a security protocol is presented where, sometimes, authentication simply cannot be achieved (due to physical constraints). When applied though, [6] also has a minimum defined tag size of 32 bits. The bandwidth-critical application of [6] is IP-packet based conversational multimedia so even if by chance a tag is forged, effects are not catastrophic. However, it is important to realize that even changing a “no” to a “yes”, may be of importance. A short tag should be seen more as a warning mechanism than as a protection mechanism.

An example where even the authentication key may be relatively short is in application of so-called Manual Authentication (MANA) Protocols. A typical application is a one-time authenticated transfer of a data item (e.g. a public key certificate) between two devices using short-range radio. A (one-time) PIN is chosen as a key and is entered (using physical side-channel) in both devices. If implemented correctly, the possibility to forge can be explicitly evaluated, and the key can sometimes be allowed to be shorter than the tag, e.g. a 16-bit key and a 160-bit tag. Note though that these arguments break down if the same key is re-used. See [35] for further details. It should also be noted that such analysis often relies on assumptions such as “physical proximity”, which is (partly) used to rule out certain attacks.

In general, when considering minimum acceptable tag size, it is important to analyze whether or not the application is designed in such a way that a possible attacker may have oracle access to a MAC verification mechanism. Short tags should not be used if this is the case.

7.1.3 User Authentication

For user authentication (using symmetric key challenge-response type protocols) the importance of the length of the response depends on the time period of validity of the authentication (the “session” length). If an attacker is lucky in guessing one response, he will most likely be detected during re-authentication (if such occurs). The impact of impersonation during a certain time period is of course application dependent, but is generally larger than the impact of faking occasional messages. Responses of 64 bits seem a general minimum. A one-time password (OTP) generator typically produces 24-32 bit values, but is usually used in connection with a longer term, user selected key (a password). Unless physical security of the communication channel is present (security against e.g. hijacking), it is important that the user authentication is coupled with a key agreement followed by integrity protection. That is, at the same time as the response is generated, a session key for integrity protection is generated from the same challenge and used to integrity protect every subsequent message. With proper replay protection of challenges, this limits the effect of a faked response to be essentially a short-lived Denial-of-Service attack.

7.1.4 Hash Functions

The main consideration is the size of the hash outputs. If the application requires collisions to be difficult to find, the output must be twice the desired security level. This is the case e.g. when used with digital signatures. When used as a keyed hash for message authentication, however, the outputs may often be truncated, see above.

7.1.5 Nonces

So-called *nonces* (number used once) are generally required to be as long as the symmetric keys used, i.e. match the security level. This is due to attacks based on the birthday paradox. For example, an n -bit nonce (here often called a salt) used for deriving a session key from a k -bit symmetric key will produce an overall security of about $(n + k)/2$ bits against off-line collision attacks.

7.2 Security Levels

Keeping in mind that the security level of Table 7.1 only gives very basic protection (a few months) we need to add some margin to get real protection. At the same time, as we just discussed, there are some applications with special properties that might allow *lower* security requirements, and there may be some very constrained environments where a certain level simply cannot be obtained. We would therefore like to define some *security levels* and quantify what security they reach and in which cases they might be acceptable. That is, rather than mandating use of certain key-sizes, we would like to be more practical and see what security is obtained in different cases.

A few comments. An 80-bit level appears to be the smallest general-purpose level as with current knowledge it protects against the most reasonable and threatening attack scenarios. The 32 and 64-bit levels should not be used for confidentiality protection; 32-bit keys offer no confidentiality at all relative to *any* attacker, and 64-bit offers only very poor protection. Nevertheless, there are applications when these levels may be necessary if security is to be provided at all, e.g. for integrity tags. While we certainly do not think this level of security should ever be promoted, it is at the same time important to be aware that some narrow-bandwidth applications would be imposed with a considerable over-head even with such short tags, and it is important to highlight what attacks would then become possible.

The choices 112/128 for the high and very high levels are a bit conservative based on extrapolation of current trends. However, since there exist well-proven standardized components that supports these levels, it still seems to make sense to define them that way.

7.3 How to Deal with Very Long-term Security

In some applications, e.g. voting, legally binding signatures, protection of state secrets, etc, very high security levels may be needed. As discussed previously, it can be difficult to obtain and maintain such high security levels. Apart from the fact that non-technical issues, outside

Table 7.4: Security levels (symmetric equivalent).

Security Level	Security (bits)	Protection	Comment
1.	32	Attacks in “real-time” by individuals	Only acceptable for auth. tag size
2.	64	Very short-term protection against small organizations	Should not be used for confidentiality in new systems
3.	72	Short-term protection against medium organizations, medium-term protection against small organizations	
4.	80	Very short-term protection against agencies, long-term prot. against small organizations	Smallest general-purpose level, < 5 years protection
5.	112	Medium-term protection	Approx. 20 years
6.	128	Long-term protection	Good, generic application-indep. recommendation, \approx 30 years
7.	256	“Foreseeable future”	Good protection against quantum computers

the scope of this report, tend to matter more, there is also a technical problem appearing. Namely, that is difficult to predict future cryptanalytic advances. As an example, a 1024-bit integer would be factored some 1000 times faster by the Number field sieve, than by the Quadratic sieve (QS) method, so yet unknown future advances could really make a difference. That is, even if extrapolation of key sizes would be a valid approach for the next, say 30 years, one cannot exclude that certain algorithms become more or less completely broken over such long time frames. So, how should we take cryptanalytic progress into account? If cryptanalysis affecting RSA had, over the last, say, 20 years an effect comparable to Moore’s law (independently, on top of the actual Moore’s law), then it would not be reasonable to assume that that progress stops at this very moment, and it could also be argued to be a stretch of the imagination that it would suddenly go much faster (even though both possibilities are in principle conceivable). For all practical purposes (i.e., a position one can defend) one should assume that future progress follows the trend that we have seen in the past and base one’s decisions on that assumption.

Advances have often been done in steps (e.g. the improvement from QS to NFS), and beyond approximately 10 years into the future, the general feeling among ECRYPT partners is that recommendations made today should be assigned a rather small confidence level, perhaps in particular for asymmetric primitives.

All is not lost though. By (again) handling security as a matter of due process, it may still be possible to maintain message authenticity and non-repudiation over long periods of

time. By keeping up to date on cryptanalytic and technological advances one could, as fear of compromise grows stronger, re-authenticate a message (and the existing signature) with larger key size and/or different algorithms.

Another approach to message authenticity is to authenticate messages with more than one algorithm. For instance, sign a message with two different algorithms, and consider the (compound) signature valid if and only if both signatures check out. It is here important that keys and parameters are generated randomly and independently, and, depending on what the feared future advances might be, it could be essential that the algorithms are very “different”. For instance, signing a message both with RSA and elliptic curve technology does not offer any additional security if quantum computers become a reality. One can also consider combinations of asymmetric and symmetric keys, e.g. both sign a message, and add symmetric key based integrity.

For confidentiality it may be more difficult to apply the security process thinking since once confidentiality is lost, it is gone forever. However, additional protection may be offered by multiple encryption, again using sufficiently “different” algorithms. The classical one-time-pad scheme may also be applicable to very high security levels, assuming the key management can be solved.

7.4 A Final Note: Key Usage Principles

Besides the above recommendations, it is important to be aware of some principles for how to use keys and other parameters.

First, the *randomness* principle must be followed: keys must be randomly or pseudo-randomly generated. If less key-material than needed is available, a pseudo-random function should be applied to obtain the required amount of key. Of course, we do not claim that the attack resistance is still any higher than the size of the input to the random number generator. Rather, we want to stress that other, ad-hoc key-expansion methods could *decrease* the security. As a basis for generating (pseudo)random bits it is crucial to properly seed one’s random number generator. This is an issue that is too often overlooked in practical applications.

Secondly, the principle of *key-separation* applies: the same key should never be used for two different purposes (e.g. use distinct keys for encryption and integrity protection, etc). Moreover, the same key should not be used twice with different transforms (e.g. encryption transform A should not use the same key as encryption transform B). For symmetric encryption, do not use the same transform inputs (key, IV, etc) twice to process different messages. For integrity protection, make sure all messages are distinct (e.g. by including a counter) to avoid replay. This recommendation holds also for other primitives, e.g. do not use the same signature key in two different systems.

Finally, for electronic signature schemes, you should not be too “promiscuous” in how you use your secret key to sign. Never sign a document provided by some other party without first applying a secure message padding scheme, including appropriate randomization, see Chapter 14.

Part II

Symmetric Primitives

Chapter 8

Block Ciphers

8.1 Overview

Block ciphers specify keyed, invertible transformations from b -bit blocks to b -bit blocks, under the influence of n -bit keys. That is, a block cipher is defined by two algorithms, E , that encrypts, and D that decrypts, such that for all n -bit keys k , and all b -bit blocks x , $D(k, E(k, x)) = x$. Most ciphers support only one block-size, but in many cases several key sizes are supported.

A number of security properties are required for a block cipher to be considered secure. Needless to say, one such requirement is that no key-recovery attack (e.g. differential [13], linear attacks [74] etc) of complexity better than 2^n is known. Stronger notions of security are for instance various notions of *indistinguishability*. What does that mean? Well, the reader may be familiar with the Turing test for artificial intelligence, where a human can “chat” (via keyboard/screen) with another entity, asking questions like “Describe in single words only the good things that come into mind about your mother”, etc. This other entity is either a human or a computer program. The program is considered artificially intelligent, if the user cannot tell if he is chatting with a human or the program. In our case, the attacker gets a plaintext, and a value which is either an encryption of said plaintext (for an unknown, random key), or, a random b -bit string. The attacker is challenged to guess which value he got, and to this end he may issue (adaptive) chosen message queries to an “encryption oracle”, on any message (except the challenge). If he cannot succeed with probability much better than one-half, the cipher is said to be indistinguishable from a random function. The intuition is that if the encipherments look just like random strings, they should not “leak” useful information about the plaintext. Various relations and equivalences are known between different security notions.

Note that a block cipher should never be used in the raw, but rather run in a specified mode of operation, e.g. cipher block chaining (CBC). CBC and many other modes require padding, and the choice of padding and implementation may affect security. In particular, error-messages related to detected padding errors may be exploited to recover messages, [116]. We stress in particular that block ciphers do not provide any real integrity protection, and without added integrity protection, also confidentiality may be lost, [11].

Also the block size may define upper bounds on the security. A small block size may enable creation of dictionaries. Also, “non-random” behaviour that might be exploitable starts to appear after about $2^{b/2}$ encrypted blocks. This chapter has therefore been divided according to block size.

For a more extensive discussion of block-cipher security properties, see the NESSIE evaluation report, [90].

8.2 64-bit Block Ciphers

8.2.1 DES

Definition: NIST FIPS 46-3, [79].

Parameters: 56-bit key and 64-bit block size

Security: Key length inadequate.

Deployment: Widespread, e.g. RFC 2406 (IPsec), RFC 2246 (TLS).

Implementation:

Public analysis: Papers such as [13, 74] (see below).

Known weakness: Susceptible to differential [13] and linear cryptanalysis [74] and their extensions, which reduces effective key size by roughly 10-12 bits. However the key-size is already acknowledged to be insufficient.

Comments: Withdrawn by NIST in 2004.

8.2.2 3DES

Definition: NIST SP-800-67, [83].

Parameters: 112-bit and 168-bit key and 64-bit block size.

Security: Taking account of the iterative construction of 3DES, attacks on 112-bit 3DES require between 2^{80} and 2^{90} operations. Similarly, attacks on 168-bit 3DES require around 2^{100} operations. In general, if 2^t plaintext/ciphertext pairs are known, two-key 3DES can be attacked with 2^{120-t} operations.

Deployment: Widespread; e.g. 112-bit 3DES widely used in financial applications, 168-bit 3DES featured within IPsec, SSL/TLS, included in ISO/IEC 18033-2, etc.

Implementation:

Public analysis: Cryptrec report [27].

Known weakness: A variety of structural properties are well-known but easily avoided.

Comments: In a draft statement 2004, 112-bit 3DES is re-affirmed by NIST for 5 years, and 168-bit 3DES is recommended to 2030.

8.2.3 Kasumi

Definition: 3GPP TS 35.202, [1].

Parameters: 128-bit key and 64-bit block size.

Security: As claimed.

Deployment: UMTS.

Implementation: 3GPP TS 35.202, [1].

Public analysis: Evaluation report [2], paper [16]. Some provable security relative to linear and differential cryptanalysis has been established [61].

Known weakness:

Comments: Variant of MISTY-1. Kasumi to be licensed for use in UMTS.

8.3 128-bit Block Ciphers

8.3.1 AES

Definition: NIST FIPS PUB 197, [82].

Parameters: 128-, 192-, and 256-bit key and 128-bit block size.

Security: As claimed.

Deployment: Widespread, included in TLS, S/MIME, IPsec, ISO/IEC 18033-3, IEEE 802.11i, etc.

Implementation:

Public analysis: NIST report [86], NESSIE and Cryptrec reports [90, 27].

Known weakness:

Comments: Algebraic attacks on the AES have been conjectured. Ongoing research will determine the likely extent of these new cryptanalytic techniques.

Chapter 9

Stream Ciphers

9.1 Overview

A stream cipher is an algorithm that takes a key (and possibly other information) and produces a (for most practical purposes) arbitrary long sequence, the *keystream*. This keystream is then combined with the cleartext to form the cipher text. In most commonly used stream cipher applications today, the keystream is bit-wise added modulo 2 to the cleartext (XORed), this is then called a *binary additive stream cipher*. Some ciphers use another combination operation, which in some cases can even provide some integrity protection. Note that it is obvious that a binary additive stream cipher in itself provides no integrity protection; an attacker can flip individual bits (even if he does not know if he flips one to zero or vice versa). Binary additive stream ciphers *should* therefore always be used with integrity protection.

Stream ciphers come in two flavours. *Synchronous* stream ciphers require sender and receiver to maintain synchronization within the key stream by external means (e.g. including an explicit synch value in each message). *Self-synchronizing* streamciphers, allow temporary loss of synch, which is then automatically regained after a while. Essentially all practically used stream ciphers today are synchronous.

It is obvious from construction that the security of a stream cipher depends on the “randomness” of the keystream, and if the stream appears indistinguishable from true random bits, security follows immediately. However, few streamciphers have managed to resist all distinguishability attacks, but if a distinguishing attack requiring huge amounts of known keystream exists, it may not be a disaster for the practical security.

Why would one choose a stream cipher rather than a block cipher? Usually, any or all of the following three properties may be motivation:

1. Stream ciphers are usually designed with high speed as a requirement, hence many stream ciphers are quite fast (though note that some constructions have at the same time suffered security-wise).
2. Stream ciphers do not require padding of messages, hence they may offer bandwidth savings in critical applications.

3. Stream ciphers do not propagate bit-errors; single transmission bit-errors in ciphertext translates to single bit-errors (in the same bit positions) after decryption.

The last motivation is probably the weakest since, as noted, it also means that attackers can flip message bits if integrity protection is not used, and integrity protection would at the same time work against the error-resistance of the application. Nevertheless, many voice-over-radio systems (without integrity protection) use stream ciphers since the voice decoders usually have high tolerance to single bit-errors, but cannot cope well with frame-errors.

Note that binary additive stream ciphers may lose all security if the same key (or more precisely, the same keystream) is used to encrypt two different messages (creating so-called two-time pad).

For a more extensive discussion of stream-cipher security properties, see the NESSIE evaluation report, [90].

9.1.1 A Note on Pseudo random Number Generation

In principle, any stream cipher can function as a pseudo random number generator (PRNG). Are there any differences in requirements on a PRNG and a stream cipher? Looking purely at the requirements on the outputs, using a strong definition of stream cipher security, there is no real difference, both need to produce outputs computationally indistinguishable from random. In practice there are usually some differences in requirements though:

- Stream ciphers are usually called for due to bandwidth saving (avoid padding) or processing speed, otherwise one can just use a block-cipher in a suitable mode (which could be a stream cipher emulating mode). A PRNG can on the other hand often be acceptable even if the speed is quite low.
- As mentioned, a distinguishing attack on a stream cipher may not be catastrophic, but if a PRNG is used to generate keys for other ciphers, it could be a completely different story. There *may* thus be reason to have stronger security requirements on a PRNG.
- Re-keying of a stream cipher is usually done to avoid “wearing out” the key or to achieve synchronization, and the requirements on avoiding key-stream re-use are quite well understood. A PRNG may be “re-seeded” to improve the randomness of the internal state, i.e. to “add entropy” to an existing state, and this is currently done more as an art than a science.

Due to the non-random properties known to exist in RC4 (the currently only included stream cipher, see below), ECRYPT recommends against using RC4 as a general PRNG, a secure block cipher in suitable mode of operation would be preferred. Future revisions of this report may cover PRNGs in more detail.

9.2 RC4

Definition: see [49, 50].

Parameters: variable key size.

Security: While no key recovery attack is known when used directly as a keystream generator, RC4 is highly sensitive to attacks which exploit re-keying and re-initialisation.

Deployment: Widespread, e.g. SSL/TLS, IEEE 802.11b, etc.

Implementation:

Public analysis: Cryptrec [27], see also e.g. [49].

Known weakness: Various distinguishing attacks apply, some key recovery attacks are known for specific implementations with re-keying. The first bytes of generated keystream are particularly vulnerable to cryptanalysis.

Comments: Recommendations often include dropping the first 512 bytes of generated keystream.

Chapter 10

Hash Functions

10.1 Overview

Cryptographic hash functions are widely used in computer and network security applications. They operate on a message of (for practical purposes) arbitrary size and produce a fixed-size fingerprint, or digest, of the message. Depending on the application, some or all of a number of security properties are required from the hash function $h(\cdot)$. Potential properties that we might appeal to are:

Pre-image resistance: Given an output $y = h(x)$ (but not a corresponding input x) it is practically infeasible to find x . Such a property might be useful when one wishes to commit to the value x at some point in time while keeping the value of x secret until later.

2nd pre-image resistance: Given an output $y = h(x)$ and a corresponding input x it is practically infeasible to find another input $z \neq x$ such that $h(z) = h(x)$. This property might be used to prevent some party from changing from a committed value.

Collision resistance: It is practically infeasible to find *any* pair of distinct inputs x and z such that $h(x) = h(z)$. This property is often required to protect electronic signature schemes against forgeries. In such schemes the hash of a message is typically signed as a representation of that message. Thus if an attacker can find two inputs x and z that collide with respect to some hash function, then the attacker might be able to re-use a legitimate signature on x as a correct, but falsely-obtained, signature on z .

Random oracle property: The function $h(\cdot)$ “behaves” as a randomly chosen function. Assuming this property holds sometimes makes it possible to formally prove the security of public key encryption and signature schemes. We shall discuss this more in Section 12.1.

While one can construct contrived examples of functions that are collision resistant but not pre-image resistant, the hash function properties have been ordered in terms of the difficulty faced by an opponent, with the task of finding a pre-image being the hardest. In the absence of any analytic weaknesses, only brute force methods are available to the attacker. More

precisely, if n is the size of the hash outputs, one would from a secure hash expect around 2^n operations to be required to break the first two properties (though this decreases as the number of available targets increase) and around $2^{n/2}$ operations to break the property of collision resistance (due to the birthday paradox). Consequently, one needs to choose a secure h with a large enough n so that these numbers meet application-dependent requirements on “practical infeasibility”.

It is quite clear from construction that none of the hash functions here can be considered “random oracles” in a generic sense, and it has in fact been established that no fixed function can have all required properties, [21]. The practical implications of the so-called random oracle model has therefore been debated, we shall return to this in the asymmetric algorithm part of this report.

For a more extensive discussion of hash function security properties, see the NESSIE evaluation report, [90].

Note that some “obvious” techniques aiming at improving/repairing the security of some hash function(s) may not have the effect one hopes. For instance, simple schemes such as concatenating the output of some commonly used hashes may not help, see [60].

10.2 MD5

Definition: RFC 1321, [97].

Parameters: 128-bit hash output, max input size $2^{64} - 1$ bits.

Security: Not collision-resistant.

Deployment: Widespread; e.g. in SSL/TLS, IPsec, etc.

Implementation: C-source code available in RFC 1321, [97].

Public analysis:

Known weakness: Collisions have been found, [117], with low computational complexity when the form of the message is restricted over 1024 bits (though the total message length may be greater and the form of the remainder is unrestricted). Collisions for (valid) public key certificates have been reported, [71]. Cryptanalytic improvements should be anticipated.

Comments: Should not be used for new applications. Existing applications may depend on other properties such as pre-image and 2nd pre-image resistance and so the implications of recent cryptanalysis need careful review. Further comments are available via the ECRYPT statement on hash functions, [29].

10.3 RIPEMD-128

Definition: see [96].

Parameters: 128-bit hash output, max input size $2^{64} - 1$ bits.

Security: As claimed; collision search requires 2^{64} iterations of the compression function. However, this is no longer adequate.

Deployment: Unknown.

Implementation: see [96].

Public analysis:

Known weakness:

Comments: While collisions on RIPEMD are reported in [117], RIPEMD is a significantly different design to RIPEMD-128. However, partial attacks on 2- or 3-round reduced-versions (out of 4 rounds) should be anticipated.

10.4 RIPEMD-160

Definition: ISO/IEC 10118-3, [54] (see also [96]).

Parameters: 160-bit hash output, max input size $2^{64} - 1$ bits.

Security: As claimed; collision search requires 2^{80} iterations of the compression function.

Deployment: Permissible algorithm in IPsec, IEEE P1363, and OpenPGP.

Implementation: see [96].

Public analysis: Cryptrec report [27].

Known weakness:

Comments: While collisions on RIPEMD have been reported, RIPEMD is a significantly different design to RIPEMD-160. Partial attacks on 2- or 3-round reduced-versions (out of 5 rounds) should be anticipated.

10.5 SHA-1

Definition: NIST FIPS 180-1 and NIST FIPS 180-2, [80].

Parameters: 160-bit hash output, max input size $2^{64} - 1$ bits.

Security: A new, yet unavailable, but credible result (see below) indicates that SHA-1 is *not* collision resistant in a strong, 2^{80} sense.

Deployment: Widespread (included in e.g. IKE/IPsec, IEEE P1363, ISO 10118-3, etc).

Implementation: RFC 3174.

Public analysis: Cryptrec report [27].

Known weakness:

Comments: While collisions on SHA have been reported, SHA is a significantly different design to SHA-1. Collisions for 36-round reduced-versions (out of 80 rounds) have been found and collisions for 53-round reduced-versions (out of 80 rounds) have been speculated.

Important notice: After finalizing a first version of this report, preliminary reports on attacks on full SHA-1, finding collisions with a 2^{69} effort have appeared in February 2005, see [102]. We believe these reports to be highly credible. We therefore recommend to be very cautious with any new deployment of SHA-1, in particular in signature schemes, since the nature/severity of the collisions that may be obtained is yet unknown at the time of writing, and since further progress (even more efficient attacks) can not be excluded.

10.6 SHA224, SHA256

Definition: NIST FIPS 180-2, [80].

Parameters: 224-bit and 256-bit hash outputs respectively, max input size $2^{64} - 1$ bits.

Security: As claimed; collision search requires 2^{112} and 2^{128} iterations of the compression function respectively.

Deployment: Likely to become widespread. Included in e.g. ISO 10118-3.

Implementation:

Public analysis: Cryptrec report [27].

Known weakness:

Comments: Collisions on SHA have been reported. While SHA has some similarities, it is also a significantly different design to SHA-224 and SHA-256. SHA-224 is identical to SHA-256, except that it uses a different IV and truncates the output.

10.7 SHA384, SHA512

Definition: NIST FIPS 180-2, [80].

Parameters: 384-bit and 512-bit hash outputs respectively, max input size $2^{128} - 1$ bits.

Security: As claimed; collision search requires 2^{192} and 2^{256} iterations of the compression function respectively.

Deployment: Unclear. Included in e.g. ISO 10118-3.

Implementation:

Public analysis: Cryptrec report [27].

Known weakness:

Comments: Collisions on SHA have been reported. While SHA has some similarities, it is also a significantly different design to SHA-384 and SHA-512. SHA-384 is identical to SHA-512, except that it uses a different IV and truncates the output.

10.8 Whirlpool

Definition: ISO/IEC 10118-3, [54] (see also [51]).

Parameters: 512-bit hash output, max input size $2^{256} - 1$ bits.

Security: As claimed; collision search requires 2^{256} iterations of the compression function.

Deployment: Unclear.

Implementation: see [51].

Public analysis:

Known weakness:

Comments: Built using AES-like components and has a different design philosophy to the MD-family.

Chapter 11

Message Authentication Codes

11.1 Overview

MACs aim to provide integrity protection. Given a key, k , they operate on a message, M , by computing an m -bit check-value $\text{MAC}(k, M)$, which is then usually appended to M before transmission/storage. The receiver/retriever of the message similarly computes the “expected” MAC value and compares it to the presented one.

The basic security notion for MACs is that without knowledge of the key, it should be infeasible for an attacker (even after seeing many, possibly even chosen, $M, \text{MAC}(k, M)$ -pairs) to produce a new, (M', t') pair so that $t' = \text{MAC}(k, M')$, other than by pure chance, guessing the value with probability 2^{-m} . Attacks included in this notion are:

Key recovery: retrieve the n -bit key, k .

Insertion: create (M, t) such that $t = \text{MAC}(k, M)$.

Modification: observing some $(M, t) = (M, \text{MAC}(k, M))$ try to modify it into (M', t') so that $t' = \text{MAC}(k, M')$.

Note that the first attack is *verifiable*, whereas the two other usually are not: the attacker knows when he has found the right key, but may not be able to tell if a candidate tag value is correct. So, both n and m are upper bounds on the security with respect to different attacks, but a smaller m can usually be accepted since non-verifiability of “tag-guessing” would only create occasional forgeries, and/or, since many applications have a natural built-in re-try limit.

Often, integrity is provided in addition to confidentiality, and one often sees discussion on whether to use an encrypt-then-authenticate, or, an authenticate-then-encrypt strategy. From security point of view, the best way is to first encrypt, then authenticate the encrypted value. First, this has the benefit that the receiver can authenticate before spending resources decrypting. Secondly, it seems intuitively clear that adding integrity check values (a form of “redundancy”) before encrypting, cannot improve confidentiality, and will (at least in theory) actually reduce security.

For a more extensive discussion of MAC security properties, see the NESSIE evaluation report, [90]. Annex B of ISO/IEC 9797-1, [52], also provides a nice summary for the algorithms included in that standard.

11.2 HMAC

Definition: RFC 2104 [67].

Parameters: A key and a hash function, often MD5 or SHA1. HMAC-MD5 provides m -bit tag for $0 \leq m \leq 128$ HMAC-SHA-1 provides m -bit tag for $0 \leq m \leq 160$. Key size depends on the hash function and the standard.

Security: As for all iterated MAC constructions, the security is not only limited by the size of the MAC tag, but also by the square root of the size of the internal state (due to birthday attacks). For HMAC-MD5 this results in a security level of 2^{64} and for HMAC-SHA-1 a security level of 2^{80} . The HMAC construction has provable security under certain assumptions on the hash function, [7].

Deployment: Widespread: SSL/TLS, IPsec, etc. Included in ISO 9797-2, NIST FIPS 198.

Implementation:

Public analysis: NESSIE report, [90].

Known weakness:

Comments: Note that increasing the key length beyond the hash function input block size (512 bits for MD5 and SHA1) does not offer any increase in security. Some standards (e.g. IPsec) allow the hash output to be truncated. Security proofs for HMAC depend upon plausible (but untested) properties of the hash functions covered elsewhere in this algorithm report. While not immediately threatening HMAC-MD5, the recent advances in the cryptanalysis of MD5 suggest that implementers move away from HMAC-MD5 as soon as possible.

Note: Please also refer to Section 10.5 for latest developments surrounding collision resistance of SHA-1. In light of this, the same caveats as mentioned for MD5 apply.

11.3 CBC-MAC-X9.19

Definition: ANSI X9.19 [4] (a.k.a. ANSI Retail MAC).

Parameters: 112-bit key and 64-bit MAC output (with optional truncation)

Security: Beyond 2^{32} MAC operations, security breaks down allowing MAC forgery and efficient key recovery attacks, [95].

Deployment: Widespread

Implementation:

Public analysis: e.g. [95] (see above).

Known weakness: A wide-range of attacks requiring 2^{32} MAC operations are known.

Comments: Implementers are recommended to move to alternative schemes for future applications unless frequent re-keying is used.

11.4 CBC-MAC-EMAC

Definition: ISO/IEC 9797-1, [52].

Parameters: Block cipher dependent, but particularly suited to the AES. Key length matches AES key lengths and provides an m -bit tag for $0 \leq m \leq 128$. 64 bits is recommended.

Security: For MAC-forgery the birthday bound gives a work effort of 2^{64} operations and for key recovery the work effort is 2^k operations where k is the length of the user-supplied key.

Deployment: Potentially widespread.

Implementation:

Public analysis:

Known weakness:

Comments:

Part III

Asymmetric Primitives

Chapter 12

Mathematical Background

Below we give a brief overview of the mathematics behind asymmetric techniques, a more extensive discussion can be found in [90].

Asymmetric schemes are based on the assumed intractability of some mathematical problem—one hopes that “breaking” the scheme is essentially as hard as solving a (presumed) difficult mathematical problem. This somewhat contrasts symmetric schemes. Though mathematical theory exists for building-blocks of symmetric primitives, the construction of e.g. a block cipher is still more of an art than a science, relying on experience from known attacks, trying to put together building-blocks so as to avoid them. Are asymmetric schemes more secure? This is an often debated question, there is really nothing that guarantees this, since the security at the very bottom still rests on an unproven assumption. Nevertheless, there is a certain attractiveness in basing the security on some quite well-studied mathematical problem that nobody (often despite large efforts) has been able to solve.

12.1 Provable Security

In some cases, one can say a little more than just that the security is “based” on a mathematical problem. It is sometimes possible to formally prove that the security of an asymmetric scheme is “as high as” the difficulty of solving an underlying mathematical problem. What does this mean? Assume that the mathematical problem is denoted M (this may be the problem of integer factorization) and that the asymmetric scheme is denoted A . A proof could for instance (intuitively, and simplified) say that:

if scheme A can be broken with computational effort T , then problem instances of M can be solved with computational effort $f(T)$,

for some function f . Now, if $f(T)$ is less than the presumed (state-of-the-art) difficulty of solving instances of M , this would mean that an attack on A would be an (seemingly unlikely, or at least unexpected) break-through in computational theory, which gives some confidence in A 's security. We here have what we in computational complexity theory call a *reduction from solving M to breaking A* . Things are, however, slightly more complex and there are two issues we would like to point out.

First, the statement of the exemplified reduction above is made in the so-called *standard model*. That is, no further assumptions are made/needed, the statement/security proof is self-contained. In practice, such proofs may be difficult to find. There have therefore been proposed alternative *proof models* to the standard model. We here mention two.

The Random Oracle Model (ROM): In this model, some component of the asymmetric primitive, often a hash function, is assumed to behave precisely as a completely random function; a random oracle.

The Generic Group Model: In this model, one assumes that the group in which the arithmetic of the asymmetric scheme is carried out has no “exploitable” properties. E.g. the elements of the group can only be treated as abstract objects, on which one can perform arithmetic, but one cannot exploit things such as representation of group elements, etc. Some elliptic curve schemes can be proven secure in this model.

Either of these assumptions clearly makes the proof weaker in the sense that it then rests on even more unproven assumptions. Moreover, a main question is whether these models have anything to do with reality. In some cases they at least do not appear unreasonable, in other cases simple counterexamples can be given, showing that these properties do not hold. In particular, concerning ROM, it is known that there are (contrived, but still) protocols that *are* secure when implemented with a “real random oracle”, yet they become insecure no matter what concrete hash function one uses, [21]. Another way of reasoning is that a proof in any of these extended models at least is an indication of some soundness of the scheme. We will not debate this issue in more depth, we only remark that clearly, a proof in the standard model is qualitatively preferred.

Besides qualitative properties of the proofs, there is also a second, quantitative aspect. In the example above, it was stated that “If scheme A can be broken with computational effort T , then problem M can be solved with computational effort $f(T)$ ”. An important question is how $f(T)$ grows as a function of T . This is usually referred to as the *tightness of the reduction from M to A* . We prefer a function f which does not grow too fast, say e.g. $f(T) = T^2$, or even linearly; $f(T) = 7T$. Why is this? First of all, if f grew too fast, $f(T)$ might be larger than the state-of-the-art method to solve M , and the proof would say nothing¹. But even if it is not that bad, one can make another, qualitatively more accurate interpretation of the reduction as

A is at most a little easier to break than problem M is to solve.

The amount of “little easier” is what is determined by f , and hence by the tightness of the reduction. The tighter the reduction, the stronger the confidence in the security of A . We will informally in the sequel just refer to reductions as *tight* or *loose*. Exact details can be found in [90], or referenced papers.

Finally, we note that sometimes, the problem M could also be the problem of breaking a(nother) cryptosystem, i.e. one then shows that A is essentially as secure as M .

¹We have cheated by suppressing the fact that the size of the problem instances of M (which determines complexity of M) typically depends on the key size, n , of A .

12.2 Choice of Primitive

The primitives in this report (and their security proofs, where such exist) are based on reductions from the RSA/factorization problem or from the discrete logarithm problem in a suitable group. To briefly recapitulate, the factoring assumption (for our purposes) states that given $N = pq$, for primes p, q of roughly equal size, it is “hard” to find the individual factors p, q . The (possibly strictly) stronger RSA assumption states that it is hard to invert the RSA function $x^e \bmod N$, $e \geq 3$. Finally, the discrete logarithm assumption states that given g^x (in a suitable group), it is “hard” to find x . A number of additional options and considerations also arise when selecting asymmetric primitives.

For factoring-based schemes, there are two main choices. First, the basic primitive: RSA or Rabin. Rabin, which is basically RSA with fixed exponent $e = 2$ is equivalent to factoring, but general RSA is not known to be. This can be viewed as an advantage for Rabin, but it also implies that being able to break confidentiality (or forge signatures) of one message is equivalent to finding the secret key.

For discrete-log based schemes, there are two main choices. First, the choice of group (either the multiplicative group modulo a prime or an elliptic curve) and whether the group is part of the public key or is it common to all users of the scheme? Secondly, the choice of the actual scheme where different security proofs are known for different schemes, more on this below. Some DLOG based schemes actually rely on the *Computational Diffie-Hellman assumption* (CDH): given g^x, g^y , it is “hard” to find g^{xy} , which is no harder than computing discrete logarithms. Yet other schemes rely on the (possibly strictly) stronger *Decisional Diffie-Hellman Assumption* (DDH): given g^x, g^y, g^z , it is “hard” to tell if $z = xy$ or not. Some relations between the hardness of these problems are studied in [75, 76, 77].

Pairing based schemes (based on groups where there is an assumed gap between CDH and DDH), or schemes based on completely different mathematical problems, could of course also be considered. No such scheme is included due to lack of deployment and/or maturity.

12.3 Non-cryptographic Attacks

As with most public key schemes, various attacks based on timing analysis, power analysis or fault analysis may need to be considered, depending on the deployment environment, see e.g. [90].

Chapter 13

Public-key Encryption

13.1 Overview

Somewhat simplified, an asymmetric encryption scheme consists of three algorithms (G, E, D) . Here, G is the key generation algorithm that generates private/public key pairs (e, d) . Anybody knowing e can encrypt messages $c = E_e(m)$ so that only the party knowing the corresponding d can retrieve $m = D_d(c)$. Since no deterministic, public key scheme can be secure in a strong sense, in practice the situation is a bit more complex, requiring random padding of messages, etc.

A number of options and considerations arise when selecting asymmetric encryption primitive. As usual, a main choice is whether to select a factoring based or a DLOG based schemes. Some proposed schemes also exist, whose security depend on the difficulty of certain lattice problems, e.g. NTRU-Encrypt, [46]. Here, standardization efforts are about to start, and a future update of this report may include this scheme. For the time being, however, only the more wide-spread factoring and DLOG based schemes are treated.

For factoring-based schemes, there are two main choices. First, the basic primitive: RSA or Rabin. Besides what we have already discussed, Rabin with $e = 2$ is faster than RSA with $e \geq 3$, but Rabin needs some redundancy to be sure that the decrypted message is the cleartext. The current revision of this report does not include any Rabin based scheme due to the lack of wide-spread deployment. Next, one needs to select a padding scheme. For RSA, PKCS#1v1.5 has no security proof based on a reduction to a widely studied problem. OAEP, included in the new PKCS#1v2.1 does, but relies on the ROM and the reduction of the proof is loose.

The present revision of the report does not include any general purpose DLOG based encryption scheme, mainly due to lack of deployment.

13.1.1 Security Notions

A number of security notions have been proposed for asymmetric encryption schemes. Today, a strong and probably the most widely accepted measure is that of *indistinguishability under adaptive chosen ciphertext attack*. Informally, the attacker first gets access to a decryption

oracle, which decrypts chosen ciphertexts on request. Next, the attacker gets a challenge: a pair of messages, and a value which is the encryption of one of them, with the task of telling which message it corresponds to. To this end, the attacker may issue additional decryption queries (except on the challenge ciphertext), and is considered successful if it can decide which message it got with non-trivial probability (slightly better than $1/2$). This notion of security is abbreviated IND-CCA2, and unless otherwise noted, is what we below mean by “secure”.

Again, for a more extensive discussion of the security of asymmetric schemes, we refer to the NESSIE evaluation report, [90].

13.1.2 Consideration: Hybrid Encryption

The amount of data that can be efficiently protected by the public key techniques below is rather limited, and it is usually a good idea to use hybrid encryption where a key is transported by public key means, and said key is then used with symmetric techniques to protect the actual data. Special public key based key encapsulation methods (KEMs) have been defined for this purpose, and they generally offer better (tighter) security proofs than the more general-purpose PK techniques above, see Chapter 16.

13.2 RSA/Factoring Based

13.2.1 RSA PKCS#1 v1.5

Definition: RFC 3447, [59] (see also [94]).

Parameters: integer N , product of primes p, q , encryption exponent e , decryption exponent d . N, e are public.

Security: Relies on the intractability of the RSA problem (and thus integer factoring), but security *may* be less.

Deployment: Widespread (e.g. (W)TLS, S/MIME).

Implementation:

Public analysis: Cryptrec [27], papers such as those listed in “known weakness” below.

Known weakness: Bad choices of p, q exist but are (w.h.p.) avoided by random choices for p, q of roughly the same size. Small e may open up to attacks on related messages, [23, 24, 47] (e.g. similar messages, re-using same random pads). Implementations that allow extensive adaptive chosen ciphertext attacks (about one million messages for a 1024-bit key), reporting back decryption errors, can be exploited to recover messages, [14].

Comments: Due to lack of security proof and the known vulnerabilities, we recommend whenever possible to use RSA-OAEP, or preferably, hybrid encryption based on RSA-KEM. If used, we recommend $|N| \geq 1024$. We recommend, if possible, to use large,

random e , and/or to restrict use to protection of short, random messages. We recommend not to use the same keys for encryption and signatures. Multi-prime options, using more than two primes exists, are defined in [94] and key-sizes/number of factors are in this case analyzed in [68].

13.2.2 RSA-OAEP

Definition: ISO/IEC 18033-2, [55] (see also [94, 59]).

Parameters: as for PKCS#1v1.5, plus a hash function and a mask generating function (MGF).

Security: provably as secure as RSA inversion in the ROM, but with a loose reduction, [33].

Deployment: unclear, though included in several other standards, such as IEEE P1363, S/MIME, etc.

Implementation:

Public analysis: NESSIE [90], Cryptrec [27], security proof [33] (correction of [10]).

Known weakness: Bad choices of p, q exists but are (w.h.p.) avoided by random choices for p, q of roughly the same size. Implementations must not enable attackers to distinguish between error conditions arising during the decoding process, since it may open up efficient chosen ciphertext attacks, [73].

Comments: We recommend to use $|N| \geq 1024$ and to use the SHA-family of hash functions rather than MD2, MD5 (part of the PKCS#1 specification). Pending investigation of the exact impacts of recent progress in cryptanalysis of SHA-1 (see Section 10.5), somewhat cautious use of SHA-1 may also be advisable. Due to the tighter security proof and the simpler implementation, we recommend to consider RSA-KEM together with hybrid encryption as an alternative to OAEP, see Chapter 16.

Discussion

This is the new RSA PKCS#1v2.0, included also in e.g. ISO/IEC 18033-2. Security is related, in the random oracle model, to the “RSA problem” with a very loose reduction, [33]. This is because of the quadratic-time cost of the reduction from the partial-domain RSA problem and a success probability lost from the RSA problem.

From an efficiency point of view, encryption and decryption are very similar to the plain RSA cryptosystem, except two more hashings. The advantage with respect to RSA-KEM (see Section 16.2.1) is the size of the produced ciphertexts, which, in this case, is just an element in \mathbb{Z}_N^* .

Suggestion to consider for future standards development: OAEP 3-round. Instead of a 2-round Feistel network, one can make a 3-round construction. This improves the reduction, since RSA-OAEP-3R directly reduce from the “RSA problem” (still with a quadratic cost). Furthermore, bandwidth can be improved since redundancy (the zeroes in OAEP) is no longer required.

13.3 ElGamal/Discrete log based

Currently, no general purpose DLOG based scheme is included. Several schemes (e.g. ElGamal based) have been proposed, but none has yet seen any wide-spread deployment. DLOG based schemes are however included in Chapter 16 for the purpose of protecting keys.

Chapter 14

Signatures

14.1 Overview

A signature scheme is usually defined as a triplet of algorithms (K, S, V) , where K is the Key generation algorithm, S is the Signing algorithm and V is the Verification algorithm. K generates pairs (s, v) of keys for the Signing/Verification algorithm. Only the party knowing s is able to generate a valid signature on m , $\sigma(m)$, but using V and the corresponding key v (assumed to be public information), anybody can efficiently decide if a given $(m, \sigma(m))$ pair is valid. Note that some schemes, e.g. Identity based schemes, also specify a fourth function, P , that generates parameters.

To meet security requirements (more on this below) and to allow signing of more or less arbitrary long messages, a signature scheme requires a hash function, so that the signing/verification algorithms operate on a fixed-size hash of the message. The combination of signature algorithm and hash function is called a *signature suite*. As discussed earlier, the hash output should be twice the “security level” (in bits). In principle, any secure hash/signature combination could be used. However, some issues should be brought to attention. For DSS, it does not make sense to use e.g. the MD5 hash algorithm, since DSS by definition only allows the SHA-family. Even if no security issue can be seen, it would create a completely new “DSS”, which increases complexity and reduces interoperability. Secondly, a too liberal use of different hash functions may open up so-called bidding-down attacks, where the security corresponds to the *weakest* available hash function. In particular, it is important to tie the signature value to the hash function used in creating it, see [62].

Some signature schemes enable the whole message, or part of it, to be recovered from the signature. These schemes can be useful in constrained environments because only the non-recoverable part of the message need be stored or transmitted with the signature.

As for asymmetric encryption, main choices are whether to use factoring or DLOG based schemes (in the latter case also which group) and what security model/proof (if any) one finds attractive.

14.1.1 Security Notions

The today most widely used security notion for signatures is similar to that for MACs, and is called *resistance against existential forgery under adaptive chosen message attack*. Informally, this means that the attacker is allowed to have messages of his own choosing signed by a “signing oracle”, after which the attacker is to provide a single valid $(m, \sigma(m))$ -pair that he has not seen before. Below, unless otherwise noted, “secure” is used in this sense. Again, for a more extensive discussion of the security of asymmetric schemes, we refer to [90].

14.2 RSA/Factoring Based

14.2.1 RSA PKCS#1 v1.5

Definition: RFC 3447, [59] (see also [94]).

Parameters: integer N , product of primes p, q , a signing exponent d , verification exponent e where N, e are public, and a hash function.

Security: Relies on the intractability of the RSA problem (and thus integer factoring), but security *may* be less.

Deployment: Widespread (e.g. (W)TLS, S/MIME).

Implementation:

Public analysis: Cryptrec [27].

Known weakness: Bad choices of p, q exists but are (w.h.p.) avoided by random choices for p, q of roughly the same size. d must not be small.

Comments: We recommend to use at least 160-bit hash functions and $|N| \geq 1024$, $e > 65536$, smaller e may be used if performance is critical. However, due to the lack of security proof, we recommend whenever possible to use RSA PSS instead, there is no advantage in using v1.5. We recommend not to use the same keys for encryption and signatures, nor using the same key with both RSA PSS and v1.5. We recommend against the use of the MD5 hash function in new implementations.

Note: Pending investigation of the exact impacts of recent progress in cryptanalysis of SHA-1 (see Section 10.5), it is advisable to also avoid SHA-1 in new deployments.

14.2.2 RSA-PSS

Definition: RFC 3447, [59] (see also [94]).

Parameters: as for v1.5 plus a mask generating function (MGF).

Security: Provably as secure as the RSA problem in the ROM (tight reduction) [58].

Deployment: Proposed for inclusion in standards (e.g. IEEE P1363), but deployment/use is unclear.

Implementation:

Public analysis: NNESSIE and Cryptrec reports [90, 27], several papers, e.g. [58, 25].

Known weakness: same as for v1.5.

Comments: Recommended alternative to v1.5, but we recommend not to use the same key. We recommend to use at least 160-bit hash functions and $|N| \geq 1024$, $e > 65536$, though small e may be used if performance is critical. We recommend against the use of the MD5 hash function. RSA PSS-R, giving message recovery, is specified in ISO/IEC 9796-2.

Note: Pending investigation of the exact impacts of recent progress in cryptanalysis of SHA-1 (see Section 10.5), it is advisable to also avoid SHA-1 in new deployments.

14.3 ElGamal/Discrete Log Based

14.3.1 DSA

Definition: FIPS PUB 186-2 (part of DSS), [81].

Parameters: p , an $L = 1024$ bit prime defining a prime field, and q , a 160-bit prime divisor of $p - 1$ defining a cyclic subgroup. Parameters can be generated according to method described in standard the [81].

Security: for proper choice of parameters, the best known attack has roughly the same complexity as that of factoring L -bit numbers, or, 2^{80} , whichever is smaller.

Deployment: Widespread. (Included in numerous standards, e.g. IKE, (W)TLS, IEEE P1363, ISO/IEC 9796-3, etc.)

Implementation: Test vectors available in [81].

Public analysis: several papers, e.g. [87, 114, 115] (details below), Cryptrec [27].

Known weakness: sensitive to “misuse”, e.g. predictable nonces [87], malicious parameter generation [114], but easily avoided. An unpublished attack exploiting “skewness” of the distribution of the random number generator from FIPS 186 claims a workfactor of 2^{64} and requires 2^{22} known signatures. A remedy is suggested in FIPS 186-2.

Comments: FIPS 186-2 defines the Digital Signature Standards (DSS) which specifies both the signature algorithms as well as the use of the SHA-1 hash function. For interoperability reasons, we do not recommend the use of DSA with other hash functions. FIPS 186-3 will specify larger keys and hashes. There exists variants on DSA, e.g. G(erman)DSA, K(orean)DSA, etc. While we see no security problems with these, we recommend the use of DSA to limit the number of options.

We recommend to verify the correctness of parameters to mitigate attacks such as those mentioned above. Even if the parameters have been already been “certified”, this gives some extra protection.

Note: Pending investigation of the exact impacts of recent progress in cryptanalysis of SHA-1 (see Section 10.5), it is advisable to also avoid SHA-1 in new deployments.

14.3.2 ECDSA

Definition: ANSI X9.62, [5] (part of DSS [81]).

Parameters: a subgroup over an elliptic curve defined over a prime or binary field. The standard [81] defines 4 curves over prime fields and 8 over binary fields, the curves corresponding roughly to $L = 160, 256, 384$, and 512-bit keys. A method to generate random curves is also specified (taken from [5]).

Security: for proper choice of parameters and L -bit keys, best known attack has complexity about $2^{L/2}$. In the generic group model, better attacks can be excluded, see [19]. However, the relevance of the generic group model for ECDSA has been questioned, [111] (see also “duplicate signatures” below).

Deployment: Widespread. (Included in numerous standards, e.g. IKE, (W)TLS, IEEE P1363, ISO 14888-3, 15946-2, etc.)

Implementation: Test vectors and implementation hints available in [81].

Public analysis: As for DSA, plus the NESSIE and Cryptrec reports [90, 27] the papers [19, 88].

Known weakness: As for DSA, in particular the possibility of badly generated curves over binary fields [115]. We recommend to avoid bad parameters choices by either using the specified curves in [81], or, to use random curves according to [5].

Comments: As for DSA. Note that also SHA-256, 384 and 512 may be required to give matching security. The use of Koblitz curves over binary fields requires slightly larger keys (the best attacks are for curves over the field \mathbb{F}_{2^m} about $\sqrt{2m}$ times faster, usually not critical). Some general concerns exist about possible future attacks on curves of “special form”, or over “special fields” e.g. [118]. As a first choice, we recommend curves over prime fields. Existence of duplicate signatures: note that each message has two valid signatures, if (r, s) is a valid signature, then so is $(r, -s)$, which in a sense means that DSA EC groups cannot be considered “generic”. Finally, an “interoperability chain” for various ECDSA standards is noted in [57]: FIPS 186-2 < ANSI X9.62 < IEEE P1363-2000 < ISO 14888-3, i.e. an implementation according to FIPS 186-2 is also compliant with ANSI X9.62, etc, but the converse may not hold.

We recommend to verify the correctness of parameters to mitigate attacks such as those mentioned above. Even if the parameters have been already been “certified”, this gives some extra protection.

Note: Pending investigation of the exact impacts of recent progress in cryptanalysis of SHA-1 (see Section 10.5), it is advisable to also avoid SHA-1 in new deployments.

Chapter 15

Public Key Authentication and Identification

15.1 Overview

In this primitive category we find protocols between a *prover* and a *verifier*. The purpose of these protocols is for the prover to be able to convince the verifier that “he is who he claims to be”. First, the protocol should be *complete*: if the prover indeed is who he claims to be, he should be able to convince any verifier (who wishes to be convinced) that this is the case. Secondly, the protocol should be *sound*: nobody but the real prover should be able convince anybody of this fact. Put differently, “convincing” here means ability to prove knowledge of some secret information, that is in some correspondence to some publicly available information such as an identity/public key, available to the verifier.

At present, the only included scheme is the GQ-scheme by Guillou and Quisquater, which is a *zero knowledge identification protocol*. Intuitively, this means that no other information, except the fact that the other party is who he/she claims to be, is revealed. In particular, no information about the secret key is leaked.

Note that with any identification protocol, the identity/authenticity of the other party is only assured at the time of protocol execution. Thus, if granting access or service is based on success of the protocol, a trusted path must also be in place, or be cryptographically established by tying together the identification with key agreement and integrity protection.

Again, for a more extensive discussion of the security of asymmetric identification schemes, we refer to [90].

15.2 GQ

Definition: ISO/IEC 9798-5, [53].

Parameters:

system-wide:

- public: an RSA public key, (N, e) , where e is prime
- secret (known only to “system authority”): a secret RSA exponent, d (corresponding to e)

user-specific:

- public key: an integer G
- secret key: integer Q (such that $G \equiv Q^{-e} \pmod{N}$)

security parameter: t , the number of repetitions

Security: asymptotically, if $te < A \log^B n$ (for some constants A, B), the scheme is (perfect) zero knowledge and, if a certain (strengthened) RSA assumption holds, the forgery success of an attacker is bounded by e^{-t} , see [20, 39, 40, 9].

Deployment: unclear

Implementation:

Public analysis: NESSIE, [90].

Known weakness: as with other known ZK identification protocols, implementations in environments where it is possible to “reset” the prover (e.g. removal of power from smart-card) may fail to provide security.

Comments: The scheme can be made identity based by letting $G = f(ID)$, where ID encodes some user-identity, and f is a suitable function. A usual recommendation is to use $t = 1$ so that the choice of e directly determines the security level. Since factorization of N is unknown, attacks based on fault analysis of CRT implementations in end-user devices can be excluded. Also, some side-channel attacks can be excluded due to the use of public exponents, [90].

Chapter 16

Key Agreement and Key Distribution

16.1 Overview

The methods can be divided into key encapsulation methods (KEMs) that can be used to transport a given a key, and key establishment protocols, where a shared key is produced as a result of the protocol. In the latter case the parties involved have little or no control over the actual key value produced, whereas in the first case one party usually has full control over the resulting key.

The security notion used for the asymmetric primitive in a KEM is basically the same as for general public key encryption, though considering that the protected value is to be used as a key, a secure key derivation function is also needed. Similarly, for Diffie-Hellman key-agreement variants, security means that observing the exchanged public values should reveal no useful information about the agreed secret.

A central question in cryptography is the relation between the Diffie-Hellman problem (CDH), its variants (e.g. DDH), and the DLOG problem. Some partial results, stating assumptions/conditions under which equivalence holds/does not hold, can be found in e.g. [75, 76, 77].

As with most public key schemes, various attacks based on timing analysis, power analysis, or fault analysis may need to be considered, depending on the deployment environment, see e.g. [90].

The schemes can form the basis for hybrid encryption in which the agreed/transported key is used in a symmetric cipher to encrypt arbitrary data. The standard ISO/IEC 18033-2 [55] specifies complete algorithm suites for hybrid encryption.

16.2 Factoring Based

16.2.1 RSA-KEM

Definition: ISO/IEC 18033-2, [55].

Parameters: same as for the basic RSA scheme (N, e, d) , plus a secure key derivation function (KDF).

Security: provably as secure as RSA inversion, very tight reduction with KDF modelled in the ROM, see e.g. [105].

Deployment: unclear

Implementation:

Public analysis: NESSIE [90]. The paper [48] shows that (asymptotically) any block of $O(\log \log N)$ bits is as secure as whole message (loose reduction).

Known weakness: same caveats as for general RSA key generation, easily avoided.

Comments: We recommend $|N| \geq 1024$ and $e > 65536$, though smaller e may be used if performance is critical. Due to the homomorphic properties of RSA, security requirements on the KDF may be higher than some other schemes. We recommend following the ISO 18033-2 specification for KDFs. RSA-KEM is an improvement of RSA-REACT. Encryption can be quite efficient, while decryption requires an inversion of the RSA function. As consequence the scheme remains very similar to the plain RSA cryptosystem (from an efficiency point of view).

16.3 DLOG Based

16.3.1 PSEC-KEM

Definition: ISO/IEC 18033-2, [55].

Parameters: a cyclic group (subgroup of a finite field or elliptic curve) of size q , and a KDF.

Security: modelling the KDF in the ROM, the scheme is provably as secure as the CDH problem with a tight reduction, see [105].

Deployment: unclear

Implementation:

Public analysis: NESSIE [90].

Known weakness:

Comments: Normal security considerations for parameter choice applies. We recommend to use prime fields of at least 1024 bits or a suitable elliptic curve, in both cases with at least 160-bit q . Binary fields may be used if performance is an issue.

16.3.2 ACE-KEM

Definition: ISO/IEC 18033-2, [55].

Parameters: as for PSEC-KEM, plus a hash function

Security: under the assumption that the hash is 2nd preimage resistant and that the KDF has certain pseudorandom properties, the scheme is provably as secure as the DDH problem with a tight reduction (i.e. in the standard model), see [26].

Deployment: unclear

Implementation:

Public analysis: NESSIE [90].

Known weakness:

Comments: Normal security considerations for parameter choice applies. We recommend to use prime fields of size at least 1024 bits or a suitable elliptic curve, in both cases with at least 160-bit q and hash function. Binary fields may be used if performance is an issue. Different security proofs under different assumptions are known, see [90] for an overview. In particular it has been shown in [26] that ACE-KEM is at least as secure as ECIES-KEM (which is also included in ISO 18033-2). Implementations should not reveal cause of decryption errors. Doing so does not obviously open up for attacks, but the security proof no longer holds.

16.3.3 Diffie-Hellman Variants

Diffie-Hellman (DH) is used as part of other security protocols, and we have chosen to include Diffie-Hellman variants according to a few of the most commonly used Internet protocols.

Some protocols/applications allow for DH to be used without authentication, often referred to as *opportunistic mode*. Note that without authentication, DH is totally vulnerable to an active man-in-the-middle attack: Alice and Bob think they have exchanged a key, but in reality there has been created one key between Alice and the attacker, and one key between the attacker and Bob. Consequently, the attacker can eavesdrop on all communication, acting as relay between Alice and Bob. Therefore, DH may fail to provide any security when active attacks are possible.

DH According to the Internet Key Exchange (IKE)

Definition: RFC 2409, [43].

Parameters: The IKE protocol suggests four cyclic groups (subgroup of a finite field or elliptic curve) of different sizes for key exchange. Later, 6 more prime field groups have been added [65]. The choice of the original groups originates from the Oakley protocol [92]. A secure PRF is also required.

Security: as claimed.

Deployment: Widely used to secure Internet connections.

Implementation: NIST PlutoPlus reference implementation for Linux [85].

Public analysis: The security of Oakley groups 3 and 4 are analysed in [108] without finding any practical attack or weaknesses.

Known weakness: The Oakley groups 1, 2 and 3 are too small to provide adequate security.

Comments:

Bibliography

- [1] 3GPP TS 35.202, *Specification of the 3GPP confidentiality and integrity algorithms; Document 2: Kasumi algorithm specification*, available at <http://www.3gpp.org/ftp/Specs/html-info/35202.htm>.
- [2] 3GPP TR 33.908, *3G Security; General report on the design, specification and evaluation of 3GPP standard confidentiality and integrity algorithms*, available at <http://www.3gpp.org/ftp/Specs/html-info/33908.htm>
- [3] F. Almgren, G. Andersson, T. Granlund, L. Ivansson, and S. Ulfberg, *How we Cracked the Code Book Ciphers*, Report, 2000. Available via answers.codebook.org
- [4] ANSI X9.19-1996, *Financial Institution Retail Message Authentication*.
- [5] ANSI X9.62, *Public Key Cryptography for the Financial Services Industry: The Elliptic Curve Digital Signature Algorithm (ECDSA)*.
- [6] M. Baugher et al, *The Secure Real-time Transport Protocol (SRTP)*, RFC 3711, IETF 2004. www.ietf.org/rfc/rfc3711.txt?number=3711
- [7] M. Bellare, R. Canetti, and H. Krawczyk, *Keying hash functions for message authentication*, Proceedings Crypto 96, LNCS **1109**, Springer-Verlag, 1996. Full paper available at <http://www.cs.ucsd.edu/users/mihir/papers/hmac.html>
- [8] M. Bellare, A. Desai, E. Jorjani, and P. Rogaway, *A Concrete Treatment of Symmetric Encryption: Analysis of DES Modes of Operation*, Proceedings 38th IEEE FOCS, pp. 394–403, 1997.
- [9] M. Bellare and A. Palacio, *GQ and Schnorr identification schemes: Proofs of security against impersonation under active and concurrent attacks*, Proceedings of Crypto'02, LNCS **2442**, pp. 162–177, Springer-Verlag.
- [10] M. Bellare and P. Rogaway, *Optimal asymmetric encryption (how to encrypt with RSA)*, Proceedings of Eurocrypt'94, LNCS **950**, pp. 92–111, Springer-Verlag.
- [11] S. Bellovin, *Problem Areas for the IP Security Protocols*, Proceedings of the 6th Usenix Unix Security Symposium, pp. 1–16, 1996, available at www.research.att.com/~smb/papers/index.html
- [12] D. Bernstein, *Circuits for Integer Factorization: A Proposal*, Manuscript, Nov. 2001. Available via <http://cr.yp.to/papers.html>

- [13] E. Biham and A. Shamir, *Differential cryptanalysis of DES-like cryptosystems*, Journal of Cryptology, **4** (1991), 3–72.
- [14] D. Bleichenbacher, *Chosen ciphertext attacks against protocols based on the RSA encryption standard PKCS#1*, Proceedings of Crypto'98, LNCS **1462**, pp. 1–12, Springer-Verlag.
- [15] M. Blaze, W. Diffie, R. Rivest, B. Schneier, T. Shimomura, E. Thompson, and M. Wiener, *Minimal Key Lengths for Symmetric Ciphers to Provide Adequate Commercial Security*, Report of ad hoc panel of cryptographers and computer scientists, Jan. 1996. Available via <http://www.crypto.com/papers/>
- [16] M. Blunden and A. Escott, *Related Key Attacks on Reduced Round KASUMI*, Proceedings FSE 2001, LNCS **2355**, pp. 277–285, Springer-Verlag.
- [17] M. Boden and S. Kowalski, *Value based risk analysis: the key to succesful commercial security targets for the Telecom Industry*, Proceedings of 2nd Common Criteria Conference, 2002.
- [18] J. R. T. Brazier, *Possible NSA Decryption Capabilities*, Manuscript 1999, available via jya.com/nsa-study.htm
- [19] D. R. L. Brown, *Generic Groups, Collision Resistance, and ECDSA*, available at <http://eprint.iacr.org/2002/026/>
- [20] M. Burmester, *An almost-constant round interactive zero-knowledge proof*, Information Processing Letters, **42**:2, 81–87, 1992.
- [21] R. Canetti, O. Goldreich, and S. Halevi, *The Random Oracle Methodology, Revisited*, In Proceedings of 30th Annual ACM Symposium on the Theory of Computing, pp. 209–218, May 1998, ACM.
- [22] I. L. Chuang, N. Gershenfeld, and M. Kubinec, *Experimental Implementation of Fast Quantum Searching*, Physical Review Letters, **80**:15 (1998), 3408–3411.
- [23] D. Coppersmith, M. Franklin, J. Patarin and M. Reiter, *Low-Exponent RSA with Related Messages*, Proceedings of Eurocrypt '96, LNCS **1070**, pp. 1–9, Springer Verlag.
- [24] J.-S. Coron, M. Joye, D. Naccache and P. Paillier, *New Attacks on PKCS #1 v1.5 Encryption*, Proceedings of Eurocrypt 2000, LNCS **1807**, pp. 369–379, Springer Verlag.
- [25] J.-S. Coron, M. Joye, D. Naccache and P. Paillier, *Universal Padding Schemes for RSA*, Proceedings of Crypto 02, LNCS **2442**, pp. 226–241, Springer-Verlag.
- [26] R. Cramer and V. Shoup, *Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack*, Cryptology ePrint Archive, Report 2001/108, 2001.
- [27] Cryptrec report annual 2002, available at www.ipa.go.jp/security/enc/CRYPTREC/fy15/doc/c02e_report2.pdf
- [28] distributed.net, *Project RC5*, available via <http://www.distributed.net/rc5/>

- [29] ECRYPT NoE, *Recent Collision Attacks on Hash Functions: ECRYPT Position Paper*, ECRYPT document STVL-ERICS-2-HASH_STMT-1.0, Nov. 2004.
- [30] ECRYPT NoE, *ECRYPT Yearly Report on Standardization*, ECRYPT document ECRYPT-ERICS-2-D.SPA_9-1.0, Jan. 2005.
- [31] EFF, Website of the electronic frontier foundation. <http://www.eff.org/descracker.html>.
- [32] SR 002 176 V1.1.1 Special Report, *Electronic Signatures and Infrastructures (ESI); Algorithms and Parameters for Secure Electronic Signatures*, ETSI, March 2003.
- [33] E. Fujisaki, T. Okamoto, D. Pointcheval, and J. Stern, *RSA-OAEP is secure under the RSA assumption*, Proceedings of Crypto'01, LNCS **2139**, pp. 260–274, Springer-Verlag.
- [34] C. Gaj et al., *Fast Implementation and Fair Comparison of the Final Candidates for Advanced Encryption Standard Using Field Programmable Gate Arrays*, In *CT-RSA 2001*, LNCS **2020**, pp. 84–99.
- [35] C. Gehrman and K. Nyberg, *Security in Personal Area Networks*, In C. Mitchell (Ed.): *Security for Mobility*, IEE 2003.
- [36] J. Gilmore (Ed.), *Cracking DES: Secrets of Encryption Research, Wiretap Politics & Chip Design*, Electronic Frontier Foundation, O'Reilly & Associates, 1998.
- [37] L. Granboulan, *How to repair ESIGN*, SCN'02, also available at <http://eprint.iacr.org/2002/074/>, 2002.
- [38] L. K. Grover, *A fast quantum mechanical algorithm for database search*, Proceedings of the 28th ACM STOC, pp. 212–219, 1996.
- [39] L. C. Guillou and J.-J. Quisquater, *A “paradoxical” identity-based signature scheme resulting from zero-knowledge*, Proceedings of Crypto'88, LNCS **403**, pp. 216–231, Springer-Verlag, 1988.
- [40] L. C. Guillou and J.-J. Quisquater, *A practical Zero-Knowledge protocol fitted to security microprocessor minimizing both transmission and memory*, Proceedings of Eurocrypt'88, LNCS **330**, pp. 123–128, Springer-Verlag.
- [41] I. Hamer and P. Cho, *DES Cracking on the Transmogrifier 2a*, In Ç. K. Koç and C. Paar (Eds.), *Cryptographic Hardware and Embedded Systems, 1st International Workshop, CHES 1999 Proceedings*, LNCS **1717**, pp. 13 – 24. Springer.
- [42] H. Handschuh and B. Preneel, *Minding your MAC algorithms*, Draft 2004, Submitted to ISB Journal.
- [43] D. Harkins and D. Carrel, *The Internet Key Exchange (IKE)*, RFC 2409, IETF.
- [44] Helion, Website: <http://www.heliontech.com/>.
- [45] A. Hodjat and I. Verbauwhede, *A 21.54 gbits/s fully pipelined AES processor on FPGA*, In *Proceedings of the IEEE Symposium on Field-Programmable Custom Computing Machines*.

- [46] J. Hoffstein, J. Pipher, and J. H. Silverman, *NTRU: A Ring-Based Public Key Cryptosystem*, In *Proceedings of ANTS III*, LNCS **1423**, pp. 267–288, Springer-Verlag, 1998.
- [47] J. Håstad, *Solving Simultaneous Modular Equations of Low Degree*, SIAM J. of Computing, **17**, 336–341, 1988.
- [48] J. Håstad and M. Näslund, *The security of all RSA and discrete log bits*, J. ACM **51:2**, 187–230 (2004).
- [49] <http://www.weizmann.ac.il/~itsik/RC4/rc4.html>
- [50] <http://burtle.burtle.net/bob/rand/isaac.html>
- [51] <http://planeta.terra.com.br/informatica/paulobarreto/WhirlpoolPage.html>
- [52] ISO/IEC 9797-1:1999, *Information technology — Security techniques — Message Authentication Codes (MACs) — Part 1: Mechanisms using a block cipher*.
- [53] ISO/IEC 9798-5:1999, *Information technology — Security techniques — Entity authentication — Part 5: Mechanisms using zero knowledge techniques*.
- [54] ISO/IEC 10118-3:2003, *Information technology — Security techniques — Hash-functions — Part 3: Dedicated hash-functions*.
- [55] ISO/IEC 18033-2, *Information technology — Security techniques — Encryption algorithms — Part 2: Asymmetric Ciphers, draft, 2004*.
- [56] IEEE P1363-2000, *Standard Specification for Public-Key Cryptography*.
- [57] D. Johnson, A. Menezes, and S. Vanstone, *The Elliptic Curve Digital Signature Algorithm (ECDSA)*, Submission to NESSIE.
- [58] J. Jonsson, *Security proofs for the RSA-PSS signature schemes and its variants*, available at <http://eprint.iacr.org/2001/053/>, 2001.
- [59] J. Jonsson, B. Kaliski, *Public-Key Cryptography Standards (PKCS) #1: RSA Cryptography Specifications Version 2.1*, RFC 3447, IETF.
- [60] A. Joux, *Multicollisions in iterated hash functions, application to cascaded constructions*, Proceedings of Crypto 04, LNCS **3152**, pp. 306–316, Springer Verlag, 2004.
- [61] J. S. Kang, S. U. Shin, D. Hong, and O Yi, *Provable security of KASUMI and 3GPP encryption mode f8*, Proceedings of ASIACRYPT 2001, LNCS **2248**, pp. 255–271, Springer-Verlag.
- [62] B. Kaliski, *Hash Function Firewalls in Signature Schemes*, RSA Conference 2002, LNCS **2271**, pp. 1–16, Springer-Verlag.
- [63] B. Kaliski, *TWIRL and RSA Key Size*, Available via www.rsasecurity.com/rsalabs
- [64] T. Kerins, E. Popovici, A. Daly and W. Marnane, *Hardware encryption engines for e-commerce*, In *Proceedings of Irish Signals and Systems Conference, ISSC 2002*, pp. 89–94.

- [65] T. Kivinen and M. Kojo, *More Modular Exponential (MODP) Diffie-Hellman groups for Internet Key Exchange (IKE)* RFC 3526, IETF.
- [66] F. Koeune, G. Rouvroy, F.-X. Standaert, J.-J. Quisquater, J.-P. David and J.-D. Legat, *A FPGA Implementation of the Linear Cryptanalysis*, In *12th International Conference on Field Programmable Logic and Applications (FPL 2002)*, Montpellier, France.
- [67] H. Krawczyk, M. Bellare, and R. Canetti, *HMAC: Keyed-Hashing for Message Authentication*, IETF RFC 2104, available at <http://www.ietf.org/rfc/rfc2104.txt?number=2104>
- [68] A. K Lenstra, *Unbelievable security; matching AES security using public key systems*, Proceedings Asiacypt 2001, LNCS **2248**, pp. 67–86, Springer-Verlag 2001.
- [69] A. K. Lenstra, A. Shamir, J. Tomlinson, and E. Tromer, *Analysis of Bernstein’s factorization circuit*, Proceedings of Asiacypt 2002, LNCS **2501**, pp. 1–26, Springer-Verlag, 2002.
- [70] A. K. Lenstra and E. R. Verheul, *Selecting Cryptographic Key Sizes*, Journal of Cryptology **14**:4, 255–293, 2001.
- [71] A. K. Lenstra, X. Wang, and B. de Weger, *Colliding X.509 Certificates based on MD5-collisions*, <http://www.win.tue.nl/~bdeweger/CollidingCertificates/>
- [72] McLoone et al., *High Performance Single-Chip FPGA Rijndael Algorithm Implementations*, In *Workshop on Cryptographic Hardware and Embedded Systems — CHES 2001, Paris, France, 2001*, LNCS.
- [73] J. Manger, *A chosen ciphertext attack on RSA Optimal Asymmetric Encryption Padding (OAEP) as standardized in PKCS #1 v2.0*, Proceedings of Crypto 2001, pp. 230–238, Springer-Verlag.
- [74] M. Matsui, *Linear cryptanalysis method for DES cipher*, Proceedings of EUROCRYPT 93, LNCS **765**, pp. 386–397, Springer Verlag.
- [75] U. Maurer, *Towards the Equivalence of Breaking the Diffie-Hellman Protocol and Computing Discrete Logarithms*, Proceedings of CRYPTO ’94, LNCS **839**, pp. 271–281.
- [76] U. Maurer and S. Wolf, *Diffie-Hellman, Decision Diffie-Hellman, and Discrete Logarithms*, Proceedings of ISIT ’98, IEEE Information Theory Society, pp. 327, 1998.
- [77] U. Maurer and S. Wolf, *The Relationship Between Breaking the Diffie-Hellman Protocol and Computing Discrete Logarithms*, SIAM J. Comp., **28**:5 (1999), 1689–1721.
- [78] A. Menezes, P. C. van Oorschot, and S. Vanstone, *Handbook of Applied Cryptography*, CRC Press, 1997.
- [79] NIST, *Data encryption standard (DES)*, FIPS PUB 46-3, available at <http://csrc.nist.gov/publications/fips/fips46-3/fips46-3.pdf>
- [80] NIST, *Secure hash standard*, FIPS PUB 180-2, available at <http://csrc.nist.gov/publications/fips/fips180-2/fips180-2withchangenotice.pdf>

- [81] NIST, *Digital Signature Standard (DSS)*, FIPS PUB 186-2, Available at <http://csrc.nist.gov/publications/fips/fips186-2/fips186-2-change1.pdf>
- [82] NIST, *Advanced Encryption Standard*, FIPS PUB 197, available at <http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf>
- [83] NIST, *Recommendation for the Triple Data Encryption Algorithm (TDEA) Block Cipher*, SP 800-67, available at <http://csrc.nist.gov/publications/nistpubs/800-67/SP800-67.pdf>
- [84] NIST, *Key Management Guideline - Workshop Document*, Draft, October 2001, available via [csrc.nist.gov/encryption/kms/key-management-guideline-\(workshop\).pdf](http://csrc.nist.gov/encryption/kms/key-management-guideline-(workshop).pdf)
- [85] NIST, *PlutoPlus: An IKE Reference Implementation for Linux*, available at <http://ipsec-wit.antd.nist.gov/newipsecdoc/pluto.html>
- [86] J. Nechvatal, E. Barker, L. Bassham, W. Burr, M. Dworkin, J. Foti, and E. Roback, *Report on the Development of the Advanced Encryption Standard (AES)*, available at <http://csrc.nist.gov/CryptoToolkit/aes/round2/r2report.pdf>
- [87] P. Q. Nguyen and I. Shparlinski, *The insecurity of the digital signature algorithm with partially known nonces*, Journal of Cryptology, **15**, 151–176, 2002. Also available at <ftp://ftp.ens.fr/pub/dmi/users/pnguyen/PubDSA.ps.gz>.
- [88] P. Q. Nguyen and I. Shparlinski, *The insecurity of the elliptic curve digital signature algorithm with partially known nonces*, Design, Codes and Cryptography, 2002. Also available at <ftp://ftp.ens.fr/pub/dmi/users/pnguyen/PubECDSA.ps.gz>.
- [89] NESSIE consortium, *Portfolio of recommended cryptographic primitives*, Feb. 2003, available via <http://www.cryptonessie.org/>
- [90] NESSIE consortium, *NESSIE Security report*, available at <https://www.cosic.esat.kuleuven.ac.be/nessie/deliverables/D20-v2.pdf>
- [91] P. C. van Oorschot and M. J. Wiener, *Parallel Collision Search with Cryptanalytic Applications*, Journal of Cryptology **12**:1 (1999), 1–28.
- [92] H. Orman, *The Oakley Key Determination Protocol* RFC 2412, IETF.
- [93] H. Orman and P. Hoffman, *Determining Strengths For Public Keys Used For Exchanging Symmetric Keys*, IETF RFC 3766/BCP 86, April 2004.
- [94] RSA Labs, *PKCS# 1: RSA Cryptography Standard*, available at <http://www.rsasecurity.com/rsalabs/node.asp?id=2125>
- [95] B. Preneel and P. C. van Oorschot, *A key recovery attack on the ANSI X9.19 retail MAC*, Electronics Letters, **32**:17 (1996), 1568–1569. Available at <http://www.scs.carleton.ca/~paulv/papers/pubs.html>
- [96] RIPEMD, <http://www.esat.kuleuven.ac.be/~bosselae/ripemd160.html>
- [97] R. Rivest, *The MD5 Message-Digest Algorithm*, IETF RFC 1321, available at <http://www.ietf.org/rfc/rfc1321.txt?number=1321>

- [98] RSA Labs, *A Cost-Based Security Analysis of Symmetric and Asymmetric Key Lengths*, RSA Labs Bulletin #13, available at www.rsasecurity.com/rsalabs/
- [99] Saggese et al., *An FPGA-Based Performance Analysis of the Unrolling, Tiling, and Pipelining of the AES Algorithm*, In *FPL 2003*, LNCS **2778**, pp. 292–302.
- [100] A. Shamir, *Factoring Large Numbers with the TWINKLE Device (Extended Abstract)*, Manuscript, 2000.
- [101] A. Shamir and E. Tromer, *Factoring large numbers with the TWIRL device*, Proceedings of Crypto 2003, LNCS **2729**, pp. 1–26, Springer-Verlag, 2003.
- [102] SHA-1 announcement, available via www.schneier.com/blog/archives/2005/02/sha1_broken.html
- [103] R. Shipsey, *How long...?*, NESSIE Report NES/DOC/RHU/WP3/015/a, available via <https://www.cosic.esat.kuleuven.ac.be/nessie/reports/phase1/rhuwp3-015.pdf>
- [104] P. W. Shor, *Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer*, SIAM J. Sci. Statist. Comput., **26** (1997).
- [105] V. Shoup, *A proposal for an ISO standard for public key encryption*, Cryptology ePrint Archive, Report 2001/112, 2001. <http://eprint.iacr.org/>
- [106] R. D. Silverman, *A Cost-Based Security Analysis of Symmetric and Asymmetric Key Lengths*, RSA Laboratories Bulletin #13, April 2000.
- [107] E. Skoudis and L. Zeltser, *Malware: Fighting Malicious Code*, Prentice Hall, 2003.
- [108] N. P. Smart, *How Secure are elliptic curves over composite extension fields?*, Proceedings of EuroCrypt '01, LNCS **2045**, pp. 30–39. Springer-Verlag, May 2001.
- [109] F. X. Standaert, *Secure and Efficient Use of Reconfigurable Hardware Devices in Symmetric Cryptography*, Ph. D. thesis, Faculté des sciences appliquées, Université catholique de Louvain.
- [110] Standaert et al., *Efficient Implementation of Rijndael Encryption in Reconfigurable Hardware: Improvements and Design Tradeoffs*, In *Workshop on Cryptographic Hardware and Embedded Systems — CHES 2003, Cologne, Germany, 2003*, LNCS **2779**, pp. 334–350.
- [111] J. Stern, D. Pointcheval, J. Malone-Lee, and N. P. Smart, *Flaws in Applying Proof Methodologies to Signature Schemes*, Proceedings Crypto 2002, LNCS **2442**, pp. 93–110, Springer-Verlag. Also available at <http://www.di.ens.fr/~pointche/pub.php?reference=MaPoSmSt02>
- [112] I. Tuomi, *The Lives and Death of Moore's Law*, Available via http://www.firstmonday.dk/issues/issue7_11/tuomi/
- [113] L. M. K. Vandersypen, M. Steffen, G. Breyta, C. S. Yannoni, M. H. Sherwood, and I. L. Chuang, *Experimental realization of Shor's quantum factoring algorithm using nuclear magnetic resonance*, Nature 414 (2001), 883–887.

- [114] S. Vaudenay, *Hidden collisions on DSS*, Proceedings of Crypto'96, LNCS **1109**, pp. 83–88, Springer-Verlag, 1996.
- [115] S. Vaudenay, *The Security of DSA and ECDSA*, Proceedings of PKC'03, LNCS **2567**, pp. 309–323 Springer-Verlag, 2003.
- [116] S. Vaudenay, *Security Flaws Induced by CBC Padding — Applications to SSL, IPSEC,...* Proceedings of Eurocrypt'02, LNCS **2332**, pp. 534–545, Springer-Verlag, 2002.
- [117] X. Wang and D. Feng, X. Lai, and H. Yu, *Collisions for Hash Functions MD4, MD5, HAVAL-128 and RIPEMD*, available at <http://eprint.iacr.org/2004/199/>
- [118] Weil descent page, http://www.cs.bris.ac.uk/~nigel/weil_descent.html
- [119] B. Weis, *The Use of RSA Signatures within ESP and AH*, IETF draft <http://www.ietf.org/internet-drafts/draft-ietf-msec-ipsec-signatures-03.txt>, Nov 2004.
- [120] M. J. Wiener, *Performance Comparison of Public-Key Cryptosystems*, RSA Crypto-Bytes **4:1** (1998), 1–5.
- [121] L. C. Williams, *A Discussion of the Importance of Key Length in Symmetric and Asymmetric Cryptography*, Available via <http://www.giac.org/practical/gsec/Lorraine-Williams-GSEC.pdf>
- [122] Xilinx. Xilinx Press Release #03142. Available at http://www.xilinx.com/prs_rls/silicon_spart/03142s3_pricing.htm.

Appendix A

Glossary

Abbreviations, most of which are explained in more detail in the documents.

3GPP	3rd Generation Partnership Project
AES	Advanced Encryption Standard
ANSI	American National Standards Institute
ASIC	Application-Specific Integrated Circuit
CCA	Chosen Ciphertext Attack
CDH	Computational Diffie-Hellman Assumption
CMA	Chosen Message Attack
CPU	Central Processing Unit
CRT	Chinese Remainder Theorem
DDH	Decisional Diffie-Hellman Assumption
DES	Data Encryption Standard
DH	Diffie-Hellman
DLOG	Discrete Logarithm
DSA	Digital Signature Algorithm
DSS	Digital Signature Standard
FPGA	Field Programmable Gate Array
EC	Elliptic Curve
ECC	Elliptic Curve Cryptography
EFF	Electronic Frontier Foundation
ETSI	European Telecommunications Standards Institute
FIPS	Federal Information Processing Standard
HW	Hardware

IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Taskforce
IEC	International Electrotechnical Commission
IKE	Internet Key Exchange
ISO	International Standardization Organization
IP	Internet Protocol
IV	Initialization Value
KEM	Key Encapsulation Method
KDF	Key Derivation Function
MAC	Message Authentication Code
MD	Message Digest
MIME	Multipurpose Internet Mail Extensions
MIPS	Mega/Million Instructions Per Second
NESSIE	New European Schemes for Signatures, Integrity and Encryption
NFS	Number Field Sieve
NIST	National Institute of Standards and Technology
NIST SP	NIST Special Publication
OAEP	Optimal Asymmetric Encryption Padding
PK	Public Key
PKCS	Public Key Cryptography Standard
PRNG	Pseudo-random Number Generator
PSS	Probabilistic Signature Scheme
QS	Quadratic Sieve
RFC	Request For Comments (see www.ietf.org)
ROM	Random Oracle Model
RSA	Rivest-Shamir-Adleman cryptosystem
SHA	Secure Hash Algorithm
TLS	Transport Layer Security
UMTS	Universal Mobile Telecommunication System
WTLS	Wireless TLS
ZK	Zero Knowledge