Preface

It is possible to build a cabin with no foundations,
but not a lasting building.
Eng. Isidor Goldreich (1906–1995)

Cryptography is concerned with the construction of schemes that withstand any abuse. Such schemes are constructed so as to maintain a desired functionality, even under malicious attempts aimed at making them deviate from their prescribed functionality.

The design of cryptographic schemes is a very difficult task. One cannot rely on intuitions regarding the typical state of the environment in which the system operates. For sure, the *adversary* attacking the system will try to manipulate the environment into untypical states. Nor can one be content with countermeasures designed to withstand specific attacks because the adversary (which acts after the design of the system is completed) will try to attack the schemes in ways that are typically different from the ones envisioned by the designer. The validity of the foregoing assertions seems self-evident; still, some people hope that in practice, ignoring these tautologies will not result in actual damage. Experience shows that these hopes rarely come true; cryptographic schemes based on make-believe are broken, typically sooner than later.

In view of these assertions, we believe that it makes little sense to make assumptions regarding the specific *strategy* that the adversary may use. The only assumptions that can be justified refer to the computational *abilities* of the adversary. Furthermore, it is our opinion that the design of cryptographic systems has to be based on *firm foundations,* whereas ad hoc approaches and heuristics are a very dangerous way to go. A heuristic may make sense when the designer has a very good idea about the environment in which a scheme is to operate, yet a cryptographic scheme has to operate in a maliciously selected environment that typically transcends the designer's view.

This work is aimed at presenting firm foundations for cryptography. The foundations of cryptography are the paradigms, approaches, and techniques used to conceptualize, define, and provide solutions to natural "security concerns." We will present some of these paradigms, approaches, and techniques, as well as some of the fundamental results

obtained using them. Our emphasis is on the clarification of fundamental concepts and on demonstrating the feasibility of solving several central cryptographic problems.

Solving a cryptographic problem (or addressing a security concern) is a two-stage process consisting of a *definitional stage* and a *constructive stage*. First, in the definitional stage, the functionality underlying the natural concern is to be identified, and an adequate cryptographic problem has to be defined. Trying to list all undesired situations is infeasible and prone to error. Instead, one should define the functionality in terms of operation in an imaginary ideal model, and require a candidate solution to emulate this operation in the real, clearly defined model (which specifies the adversary's abilities). Once the definitional stage is completed, one proceeds to construct a system that satisfies the definition. Such a construction may use some simpler tools, and its security is proven relying on the features of these tools. In practice, of course, such a scheme may also need to satisfy some *specific* efficiency requirements.

This work focuses on several archetypical cryptographic problems (e.g., encryption and signature schemes) and on several central tools (e.g., computational difficulty, pseudorandomness, and zero-knowledge proofs). For each of these problems (resp., tools), we start by presenting the natural concern underlying it (resp., its intuitive objective), then define the problem (resp., tool), and finally demonstrate that the problem may be solved (resp., the tool can be constructed). In the last step, our focus is on demonstrating the feasibility of solving the problem, not on providing a practical solution. As a secondary concern, we typically discuss the level of practicality (or impracticality) of the given (or known) solution.

Computational Difficulty

The specific constructs mentioned earlier (as well as most constructs in this area) can exist only if some sort of computational hardness exists. Specifically, all these problems and tools require (either explicitly or implicitly) the ability to generate instances of hard problems. Such ability is captured in the definition of one-way functions (see further discussion in Section 2.1). Thus, one-way functions are the very minimum needed for doing most sorts of cryptography. As we shall see, one-way functions actually suffice for doing much of cryptography (and the rest can be done by augmentations and extensions of the assumption that one-way functions exist).

Our current state of understanding of efficient computation does not allow us to prove that one-way functions exist. In particular, the existence of one-way functions implies that $\mathcal{NP}$ is not contained in $\mathcal{BPP} \supseteq \mathcal{P}$ (not even "on the average"), which would resolve the most famous open problem of computer science. Thus, we have no choice (at this stage of history) but to assume that one-way functions exist. As justification for this assumption, we may only offer the combined beliefs of hundreds (or thousands) of researchers. Furthermore, these beliefs concern a simply stated assumption, and their validity follows from several widely believed conjectures that are central to various fields (e.g., the conjecture that factoring integers is hard is central to computational number theory).

Since we need assumptions anyhow, why not just assume what we want (i.e., the existence of a solution to some natural cryptographic problem)? Well, first we need

to know what we want: As stated earlier, we must first clarify what exactly we want; that is, we must go through the typically complex definitional stage. But once this stage is completed, can we just assume that the definition derived can be met? Not really. Once a definition is derived, how can we know that it can be met at all? The way to demonstrate that a definition is viable (and so the intuitive security concern can be satisfied at all) is to construct a solution based on a *better-understood* assumption (i.e., one that is more common and widely believed). For example, looking at the definition of zero-knowledge proofs, it is not a priori clear that such proofs exist at all (in a non-trivial sense). The non-triviality of the notion was first demonstrated by presenting a zero-knowledge proof system for statements regarding Quadratic Residuosity that are believed to be hard to verify (without extra information). Furthermore, contrary to prior beliefs, it was later shown that the existence of one-way functions implies that any NP-statement can be proven in zero-knowledge. Thus, facts that were not at all known to hold (and were even believed to be false) were shown to hold by reduction to widely believed assumptions (without which most of modern cryptography collapses anyhow). To summarize, not all assumptions are equal, and so reducing a complex, new, and doubtful assumption to a widely believed simple (or even merely simpler) assumption is of great value. Furthermore, reducing the solution of a new task to the assumed security of a well-known primitive typically means providing a construction that, using the known primitive, solves the new task. This means that we not only know (or assume) that the new task is solvable but also have a solution based on a primitive that, being well known, typically has several candidate implementations.

Structure and Prerequisites

Our aim is to present the basic concepts, techniques, and results in cryptography. As stated earlier, our emphasis is on the clarification of fundamental concepts and the relationship among them. This is done in a way independent of the particularities of some popular number-theoretic examples. These particular examples played a central role in the development of the field and still offer the most practical implementations of all cryptographic primitives, but this does not mean that the presentation has to be linked to them. On the contrary, we believe that concepts are best clarified when presented at an abstract level, decoupled from specific implementations. Thus, the most relevant background for this work is provided by basic knowledge of algorithms (including randomized ones), computability, and elementary probability theory. Background on (computational) number theory, which is required for specific implementations of certain constructs, is not really required here (yet a short appendix presenting the most relevant facts is included in the first volume so as to support the few examples of implementations presented here).

Organization of the Work. This work is organized in two parts (see Figure 0.1): *Basic Tools* and *Basic Applications*. The first volume (i.e., [108]) contains an introductory chapter as well as the first part (Basic Tools), which consists of chapters on computational difficulty (one-way functions), pseudorandomness, and zero-knowledge proofs. These basic tools are used for the Basic Applications of the second part (i.e., the current

Volume 1: Introduction and Basic Tools
Chapter 1: Introduction
Chapter 2: Computational Difficulty (One-Way Functions)
Chapter 3: Pseudorandom Generators
Chapter 4: Zero-Knowledge Proof Systems
Volume 2: Basic Applications
Chapter 5: Encryption Schemes
Chapter 6: Digital Signatures and Message Authentication
Chapter 7: General Cryptographic Protocols

Figure 0.1: Organization of this work.

volume), which consists of chapters on Encryption Schemes, Digital Signatures and Message Authentication, and General Cryptographic Protocols.

The partition of the work into two parts is a logical one. Furthermore, it has offered us the advantage of publishing the first part before the completion of the second part. Originally, a third part, entitled *Beyond the Basics,* was planned. That part was to have discussed the effect of Cryptography on the rest of Computer Science (and, in particular, complexity theory), as well as to have provided a treatment of a variety of more advanced security concerns. In retrospect, we feel that the first direction is addressed in [106], whereas the second direction is more adequate for a collection of surveys.

Organization of the Current Volume. The current (second) volume consists of three chapters that treat encryption schemes, digital signatures and message authentication, and general cryptographic protocols, respectively. Also included is an appendix that provides corrections and additions to Volume 1. Figure 0.2 depicts the high-level structure of the current volume. Inasmuch as this volume is a continuation of the first (i.e., [108]), one numbering system is used for both volumes (and so the first chapter of the current volume is referred to as Chapter 5). This allows a simple referencing of sections, definitions, and theorems that appear in the first volume (e.g., Section 1.3 presents the computational model used throughout the entire work). The only exception to this rule is the use of different bibliographies (and consequently a different numbering of bibliographic entries) in the two volumes.

Historical notes, suggestions for further reading, some open problems, and some exercises are provided at the end of each chapter. The exercises are *mostly* designed to help and test the basic understanding of the main text, not to test or inspire creativity. The open problems are fairly well known; still, we recommend a check on their current status (e.g., in our updated notices web site).

Web Site for Notices Regarding This Work. We intend to maintain a web site listing corrections of various types. The location of the site is

`http://www.wisdom.weizmann.ac.il/~oded/foc-book.html`

Chapter 5: Encryption Schemes
- The Basic Setting (Sec. 5.1)
- Definitions of Security (Sec. 5.2)
- Constructions of Secure Encryption Schemes (Sec. 5.3)
- Advanced Material (Secs. 5.4 and 5.5.1–5.5.3)

Chapter 6: Digital Signatures and Message Authentication
- The Setting and Definitional Issues (Sec. 6.1)
- Length-Restricted Signature Scheme (Sec. 6.2)
- Basic Constructions (Secs. 6.3 and 6.4)
- Advanced Material (Secs. 6.5 and 6.6.1–6.6.3)

Chapter 7: General Cryptographic Protocols
- Overview (Sec. 7.1)
- Advanced Material (all the rest):
  - The Two-Party Case (Sec. 7.2–7.4)
  - The Multi-Party Case (Sec. 7.5 and 7.6)

Appendix C: Corrections and Additions to Volume 1

Bibliography and Index

Figure 0.2: Rough organization of this volume.

Using This Work

This work is intended to serve as both a textbook and a reference text. That is, it is aimed at serving both the beginner and the expert. In order to achieve this aim, the presentation of the basic material is very detailed so as to allow a typical undergraduate in Computer Science to follow it. An advanced student (and certainly an expert) will find the pace (in these parts) far too slow. However, an attempt was made to allow the latter reader to easily skip details obvious to him/her. In particular, proofs are typically presented in a modular way. We start with a high-level sketch of the main ideas and only later pass to the technical details. Passage from high-level descriptions to lower-level details is typically marked by phrases such as "details follow."

> In a few places, we provide straightforward but tedious details in indented paragraphs such as this one. In some other (even fewer) places, such paragraphs provide technical proofs of claims that are of marginal relevance to the topic of the work.

More advanced material is typically presented at a faster pace and with fewer details. Thus, we hope that the attempt to satisfy a wide range of readers will not harm any of them.

Teaching. The material presented in this work, on the one hand, is way beyond what one may want to cover in a course and, on the other hand, falls very short of what one may want to know about Cryptography in general. To assist these conflicting needs, we make a distinction between *basic* and *advanced* material and provide suggestions for further reading (in the last section of each chapter). In particular, sections, subsections, and subsubsections marked by an asterisk (*) are intended for advanced reading.

Depending on the class, each lecture consists of 50–90 minutes. Lectures 1–15 are covered by the first volume. Lectures 16–28 are covered by the current (second) volume.

Lecture 1: Introduction, Background, etc. (depending on class)

Lectures 2–5: *Computational Difficulty* (*One-Way Functions*)
Main: Definition (Sec. 2.2), Hard-Core Predicates (Sec. 2.5)
Optional: Weak Implies Strong (Sec. 2.3), and Secs. 2.4.2–2.4.4

Lectures 6–10: *Pseudorandom Generators*
Main: Definitional Issues and a Construction (Secs. 3.2–3.4)
Optional: Pseudorandom Functions (Sec. 3.6)

Lectures 11–15: *Zero-Knowledge Proofs*
Main: Some Definitions and a Construction (Secs. 4.2.1, 4.3.1, 4.4.1–4.4.3)
Optional: Secs. 4.2.2, 4.3.2, 4.3.3–4.3.4, 4.4.4

Lectures 16–20: *Encryption Schemes*
Main: Definitions and Constructions (Secs. 5.1, 5.2.1–5.2.4, 5.3.2–5.3.4)
Optional: Beyond Passive Notions of Security (Overview, Sec. 5.4.1)

Lectures 21–24: *Signature Schemes*
Definitions and Constructions (Secs. 6.1, 6.2.1–6.2.2, 6.3.1.1, 6.4.1–6.4.2)

Lectures 25–28: *General Cryptographic Protocols*
The Definitional Approach and a General Construction (Overview, Sec. 7.1).

Figure 0.3: Plan for one-semester course on Foundations of Cryptography.

This work is intended to provide all material required for a course on Foundations of Cryptography. For a one-semester course, the teacher will definitely need to skip all advanced material (marked by an asterisk) and perhaps even some basic material; see the suggestions in Figure 0.3. Depending on the class, this should allow coverage of the basic material at a reasonable level (i.e., all material marked as "main" and some of the "optional"). This work can also serve as a textbook for a two-semester course. In such a course, one should be able to cover the entire basic material suggested in Figure 0.3, and even some of the advanced material.

Practice. The aim of this work is to provide sound theoretical foundations for cryptography. As argued earlier, such foundations are necessary for any *sound* practice of cryptography. Indeed, practice requires more than theoretical foundations, whereas the current work makes no attempt to provide anything beyond the latter. However, given a sound foundation, one can learn and evaluate various practical suggestions that appear elsewhere (e.g., in [149]). On the other hand, lack of sound foundations results in an inability to critically evaluate practical suggestions, which in turn leads to unsound

decisions. Nothing could be more harmful to the design of schemes that need to withstand adversarial attacks than misconceptions about such attacks.

Relationship to Another Book by the Author

A frequently asked question refers to the relationship of the current work to my text *Modern Cryptography, Probabilistic Proofs and Pseudorandomness* [106]. That text consists of three brief introductions to the related topics in its title. Specifically, Chapter 1 of [106] provides a brief (i.e., 30-page) summary of the current work. The other two chapters of [106] provide a wider perspective on two topics mentioned in the current work (i.e., Probabilistic Proofs and Pseudorandomness). Further comments on the latter aspect are provided in the relevant chapters of the first volume of the current work (i.e., [108]).

A Comment Regarding the Current Volume

There are no privileges without duties.
Adv. Klara Goldreich-Ingwer (1912–2004)

Writing the first volume was fun. In comparison to the current volume, the definitions, constructions, and proofs in the first volume were relatively simple and easy to write. Furthermore, in most cases, the presentation could safely follow existing texts. Consequently, the writing effort was confined to reorganizing the material, revising existing texts, and augmenting them with additional explanations and motivations.

Things were quite different with respect to the current volume. Even the simplest notions defined in the current volume are more complex than most notions treated in the first volume (e.g., contrast secure encryption with one-way functions or secure protocols with zero-knowledge proofs). Consequently, the definitions are more complex, and many of the constructions and proofs are more complex. Furthermore, in most cases, the presentation could not follow existing texts. Indeed, most effort had to be (and was) devoted to the actual design of constructions and proofs, which were only inspired by existing texts.

The mere fact that writing this volume required so much effort may imply that this volume will be very valuable: Even experts may be happy to be spared the hardship of trying to understand this material based on the original research manuscripts.