Review of
Foundations of Cryptography II: Basic Applications*

Riccardo Pucella

Department of Computer Science
Cornell University

July 5, 2005

This volume is the second in a series that aims at elucidating the foundations of cryptography. The first volume, *Foundations of Cryptography: Basic Tools*, emphasized the basic computational tools needed to study cryptography; this volume applies these tools to the design of practical schemes for encryption, digital signatures, and general cryptographic protocols.

Throughout both volumes, Goldreich advocates a particular approach to reasoning about security, characterized by the following very plausible starting point: we can make no assumption as to the strategies of adversaries. In fact, the only assumption that can be justified is an assumption on the computational abilities of adversaries. (For instance, the adversary may only be able to perform probabilistic polynomial-time computations.) This leads to very natural models in terms of computational complexity. From this point of view, many cryptographic constructs can exist only if some hard problems exist. This makes the existence of cryptographic constructs dependent on complexity theoretic assumptions, such as the existence of one-way functions. I will return to this point shortly.

The main technique for reasoning about security in this setting is the *simulation paradigm*: a scheme is secure if whatever a feasible adversary can obtain after attacking it is also feasibly attainable in an "ideal setting". For example, an encryption scheme is secure if whatever information an adversary can obtain after eavesdropping an encrypted message on an insecure channel, he can already obtain by eavesdropping on the communication performed on a private channel; in other words, he may as well guess what the message is.

Overview of the First Volume. *Foundations of Cryptography: Basic Tools* has been reviewed twice in SIGACT News (in vol 34, no 4, and vol 36, no 2), so a quick overview will suffice. Such an overview is probably needed, because the first volume is a prerequisite to the second; the task here is to highlight which parts of the first volume are required for the second volume.

Chapter 1 introduces the philosophy of the approach, as well as reviewing much of the background required for both volumes, namely probability theory, and computational complexity models. Chapter 2 introduces the core construct of much of cryptography, *one-way functions*: a one-way function is a function that is efficient to compute in one direction, but difficult to invert. The existence of one-way functions requires the existence of hard (on average) problems of a specific form; it should be possible to generate hard instances of those problems, along with enough "auxiliary" information to help solve the instance of the problem if one is given that auxiliary information. One-way functions are used extensively in the second volume. Chapter 3 introduces pseudorandom generators: a *pseudorandom generator* is an efficient (polynomial-time) deterministic algorithm that transforms a short randomly chosen string into a much longer "pseudorandom" string. A pseudorandom string is a string that is *computationally indistinguishable* from a true random string, that is, indistinguishable using efficient algorithms. Computational indistinguishability is used extensively in the second volume. The bulk of the chapter is devoted to the relationship between pseudorandom generators and one-way functions. Chapter 4 studies one of the most remarkable devices in complexity theory, zero-knowledge interactive proof systems. The setting is one in which one party A proves an assertion to another party

*O. Goldreich, *Foundations of Cryptography II: Basic Applications*, Cambridge University Press, 2004, 798pp, ISBN 0521830842.

B, so that B is convinced of the validity of the assertion. Thus, an interactive proof system consists of two parties, a "prover" whose task is to produce a proof of an assertion, and a "verifier" whose task is to verify the validity of a proof. In a zero-knowledge interactive proof system, the verifier B *does not learn anything* beyond the fact that the assertion is true (and its consequences). The main result of the chapter is a method for constructing zero-knowledge proofs for every language in NP: it is possible to construct zero-knowledge interactive proof systems for any given language L in NP, that proves queries about language membership in L. This construction relies intrinsically on one-way functions. Zero-knowledge proofs for languages in NP are used in Chapter 7 of the second volume.

Content of the Second Volume. *Foundations of Cryptography*'s second volume picks up where the first one left off. Literally so, since it starts with Chapter 5 on page 373.

The model followed in all of the chapters consists of dividing the design work in two stages. The first stage, the definitional stage, is where one defines formally the security property of interest. This is often nontrivial, as there are many possible reasonable definitions, and many different definitions turn out to be equivalent depending on the context. The second stage, the construction stage, is where one constructs schemes satisfying the security definition, often under some complexity-theoretic assumption, such as the existence of one-way functions. Much effort is devoted to finding constructions that rely on assumptions that are simple and widely believed to be true. Goldreich emphasizes early:

> [...] 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. (p.xv)

This goal drives much of the constructions in this second volume.

Chapter 5, "Encryption Schemes", studies classical goal of cryptography, that is, the problem of constructing schemes allowing secret data exchange over insecure channels.

- 5.1 "The Basic Setting": quickly discusses the basic setting in which encryption schemes are studied, distinguishing between private-key and public-key encryption schemes.
- 5.2 "Definitions of Security": gives the agreed-upon definition of security for encryption schemes under passive adversaries (that is, adversaries that can eavesdrop on the insecure channel, but cannot otherwise affect the behavior of the system) due to Goldwasser and Micali, semantic security. A different definition, also common, is given, indistinguishability of encryption (two different encrypted strings are computationally indistinguishable), and is shown to be equivalent to semantic security. These definitions are extended to deal with the security of a sequence of messages encrypted with the same key.
- 5.3 "Constructions of Secure Encryption Schemes": presents constructions of secure private-key and public-key encryption schemes that satisfy the definition of semantic security in the multiple-message setting, using cryptographic primitives from the first volume. Both stateful (stream-ciphers) and stateless schemes are described.
- 5.4 "Beyond Eavesdropping Security": the longest section in this chapter, describes stronger types of adversaries that can manipulate the system beyond simply eavesdropping on the insecure channel, and gives definitions of security and constructions satisfying them corresponding to the different kind of adversaries. These attacks include passive attacks with key-dependent plaintexts (where the chosen plaintext depends on the key used), chosen plaintext attacks (where the adversary may obtain encryption of plaintexts of his choice), and chosen ciphertext attacks (where the adversary may obtain decryption of ciphertexts of his choice). The relation notion of malleability is also examined, where the adversary seeks to generate encryption of plaintexts related to the secret plaintext, rather than gaining information about the latter.

Chapter 6, "Digital Signatures and Message Authentication", studies the problem of authenticating data, that is, designing methods for verifying that the data was approved by a certain party (or set of parties). Goldreich identifies

two distinct authentication mechanisms: signature schemes, which require that signatures be universally verifiable (anyone can verify a signature and be convinced of the authenticity of the data), and message-authentication schemes, which require that "authentication tags" need only be verifiable by parties that are able to generate them. Roughly speaking, the distinction is akin to the distinction between public-key and private-key encryption schemes.

- 6.1 "The Setting and Definitional Issues": discusses both forms of authentication schemes, and sets the stage for a unified treatment of the subject. This section also describes attacks and corresponding definitions of security for authentication schemes.

- 6.2 "Length-Restricted Signature Schemes": constructs basic authentication schemes used in the construction of more general schemes; these basic schemes are restricted so as to apply only to documents of a predetermined length. The idea of constructing general schemes from length-restricted ones via the important hash-and-sign paradigm (first hashing the document into a fixed-length string, and signing the string) is described, along with a discussion of collision-free hashing functions.

- 6.3 "Constructions of Message-Authentication Schemes": presents several constructions of secure message-authentication schemes, following and expanding the ideas of Section 6.2.

- 6.4 "Constructions of Signature Schemes": presents several constructions of secure signature schemes, following and expanding the ideas of Section 6.2.

- 6.5 "Some Additional Properties": discusses further properties that one may want out of an authentication scheme, such as having unique valid signatures (every document has a unique signature that corresponds to a verification key), or being super-secure (it is infeasible to find a different signature for the same verification key to a document for which a signature was obtained during an attack).

Chapter 7, "General Cryptographic Protocols", turns to the general problem of designing general (multi-party) secure cryptographic protocols. In a sense, we can view the design of an encryption scheme or a signature scheme as the design of a secure protocol for implementing a specific functionality between two parties (that is, transmitting a message secretly, or transmitting a message such that the origin of the message is validated). A functionality maps m inputs (the state of each party) to m outputs (the final state of each party) with an associated probability. The basic model here is as follows. If we assume there is a trusted third-party, any functionality can be computed by having each party send its input to the trusted third-party, who computes the outputs, and sends each party his output. The question is to what extent a functionality can be computed without a trusted third-party, when the parties taking part in the protocol do not trust each other and when some of the parties are adversaries. This chapter addresses this question in a number of contexts. It is by far the most technical, partly to deal with the extreme generality of the problem. It is in fact surprising that so many results can be obtained for so general a problem.

- 7.1 "Overview": gives an overview of the secure multi-party computation problem, and outlines the development of the chapter.

- 7.2 "The Two-Party Case: Definitions": defines security for two-party protocols, under two different adversary models: restricted adversaries called semi-honest (which follow the protocol, but keep a record of all their intermediate computations and may share them with other semi-honest adversaries), and general malicious adversaries that need not follow the protocol. In both cases, the adversaries in the system are fixed, albeit unknown.

- 7.3 "Privately Computing (Two-Party) Functionalities": constructs two-party protocols for computing functionalities that are secure in the semi-honest model.

- 7.4 "Forcing (Two-Party) Semi-Honest Behavior": shows how to construct, given a two-party protocol secure in the semi-honest model, a protocol that is secure in the malicious model. The idea is to "compile" each instruction of the protocol so as to force each party to behave in a semi-honest manner, or be detected as cheating.

– 7.5 “Extensions to the Multi-Party Case”: extends the treatment of general secure protocols to the multi-party case. The approach followed is as in the two-party case: definitions of security, constructions of protocols secure under a semi-honest adversary model, and extension of the result to malicious adversaries. Here, however, two kinds of malicious adversary models are considered: a model with an arbitrary number of malicious adversaries, and a model in which the number of malicious adversaries is strictly less than half the total number of parties.

– 7.6 “Perfect Security in the Private Channel Model”: presents an alternative treatment of general secure multi-party protocols. The model assumes the existence of private channels between each pair of parties. In this context, when the honest parties are in strict majority, it is possible to construct protocols that are perfectly secure without relying on intractability assumptions.

Every chapter concludes with miscellaneous topics on the subject of the chapter, historical notes, suggestions for further reading, as well as open problems and exercises. The historical notes are invaluable; the literature is so vast that a highlight of the main historical publications is needed. Personally, I quite enjoy having a list of open problems; it gives me a sense of where the development I am reading about fits in the big picture, and where I stand with respect to the bleeding edge of research. The second volume concludes with an appendix containing corrections and additions to the first volume.

Opinion. What follows is necessarily completely subjective. This book contains the most accessible overall theory of the foundations of cryptography available. It unifies the treatment found until now only in research papers, most of them rather obscure without a suitable guide. If only for that, the book is invaluable as a reference. Quoting Goldreich:

> 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.’ (p.xix)

I wholeheartedly agree. The fact that, despite being a digest of the existing literature presented in a unified fashion, the result is still technically dense is a tribute to the depth of the field, making a volume such as this one a necessity.

Two points, somewhat related, should be kept in mind when considering this book for a course on cryptography. First off, this book as a whole is fairly advanced—it requires quite a bit of mathematical maturity. Indeed, Goldreich marks the sections he considers advanced with a *; a quick look shows that half of Chapter 5 is marked as advanced (all of Section 5.4), while pretty much all of Chapter 7 is marked as advanced (all but the overview). Thus, the amount of material immediately accessible is slight. The second thing to keep in mind is that the focus is exclusively on the foundations of cryptography. It therefore implicitly assumes a general intuition for cryptography in the first place, the kind of intuition that one obtains by playing with simple (and not so simple) ciphers. Many of the subtleties in the various definitions of security may well be lost on a student without such intuition. Goldreich quite explicitly decides not to explore implemented schemes in the book:

> The reader may note that we have avoided the presentation of several popular encryption schemes. We regret to say that most of these schemes are proposed without any reference to a satisfactory notion of security. Thus, it is not surprising that we have nothing to say about the contents of such proposals. (p. 477)

While one may sympathize with this view, it has the consequence that a course based on this book might require either prior exposure to cryptography (perhaps in the form of an undergraduate course), or an introduction via some other textbook. This is not meant to be an uncontroversial statement, and probably trudges along well-defined pedagogical battle lines. In any event, a student might well be served by first reading a generic high-level textbook on cryptography, such as Stinson’s “Cryptography: Theory and Practice” or Schneier’s “Applied Cryptography” (to pick two at random from my bookshelf). These are nowhere as formal, but they are slightly more accessible. After having worked through such a book and seen examples of encryption schemes and the kind of attacks they are subject to, the case can be made that one needs good foundations.

Goldreich says that while there was an initial plan for a third volume, “Beyond the Basics”, that would discuss advanced security topics, he felt that this would be better addressed by a collection of surveys. I sincerely hope that he decides to put together such a collection of surveys, under a unified treatment paralleling the one in these two volumes, and ends up calling it *Foundations of Cryptography III: Beyond the Basics*. The work is just too good to stop there.