Preface to Special Issue on Encryption in the Bounded Storage Model

Oded Goldreich
Department of Computer Science and Applied Mathematics
Weizmann Institute of Science, Rehovot, ISRAEL
oded@wisdom.weizmann.ac.il

August 10, 2003

Most *special-issues* are triggered by some general idea (e.g., the desire to commemorate a special event or date, the realization that some area deserves special attention, etc). Typically, in these cases, a *call for papers* is posted, and the editor selects papers among those submitted. Thus, the decision to have a special issue precedes the specific choice of papers to be included in it.

The current special issue has evolved in the reverse order. First it occurred to me that the three papers included in the current issue would make a very nice special issue. Next, it occured to me that this is indeed a good time to draw attention to the *Bounded-Storage Model*. Finally, I initiated this special issue and asked the authors to submit these papers to it. Needless to say, the papers went through a regular review process.

Having started this preface on a personal note, let me continue by providing my personal perspective on some issues related to the Bounded-Storage Model. I aim to offer an unusually wide and somewhat controversial perspective, and refer the readers to the papers themselves for more conventional perspectives.

On making no assumptions. In some discussions of "information-theoretic security", one may hear the claim that in this arena "no assumptions are made" (and that this stands in contrast to "complexity-based Cryptography"). This claim is clearly false. For example, any cryptographic work assumes the ability to generate random secrets (i.e., to generate objects

that are unpredictable by the adversary).[1] In general, we always make assumptions about the (computational) *abilities of the legitimate parties*; that is, that the legitimate parties can perform certain actions. However, when one says that "no assumptions are made", one typically means assumptions about the (computational) *limitations of the adversary*; that is, that the adversary cannot perform certain actions, typically because its computational resources are bounded. A third level of assumptions refers to widely believed *conjectures*. These may be conjectures from Number Theory (e.g., the Extended Riemann Hypothesis), from Physics (e.g., the hypothesis that Quantum Mechanics provides a complete model of the physical world), or from Complexity Theory (e.g., the existence of one-way functions). We comment that, whereas it seems that the popular famous conjectures of Mathematics or Complexity Theory can be proved or disproved (albeit possibly not in this century), in principle, the conjectures of Physics can only be disproved.

Traditional work in "information-theoretic cryptography" make only assumptions of the first kind. Although the bounded-storage model may be viewed as part of "information-theoretic cryptography" (by placing bounds on the information available to the adversary), it can also be viewed as part of "complexity-based cryptography" (where the bounded complexity resource is that of space).[2] Either way, the bounded-storage model makes an assumption of the second kind; that is, an assumption referring to the (storage) limitations of the adversary. The fact that space-complexity is better understood than time-complexity, allows typical work in the bounded-storage model to obtain (positive) results without resorting to unproven conjectures (i.e., assumptions of the third kind).

On space-bounded versus time-bounded adversaries. As stated above, work in the bounded-storage model may be viewed as an application of space-complexity to cryptography. Arguably, this application of space-complexity is quite straightforward; that is, it does not refer to the actual

[1] Here and below, we actually combine two types of assumptions regarding the abilities of the legitimate parties. The first type refers to the ability of these parties to conduct certain actions (e.g., toss coins), whereas the second type refers to their ability to conduct actions in private (e.g., toss a coin without anybody else seeing or a-priori knowing the outcome).

[2] Indeed, the intimate relation between space-complexity and information theory is evident in works such as Nisan's work "Pseudorandom Generators for Space Bounded Computation" (*Combinatorica*, Vol. 12 (4), pages 449–461, 1992) and Nisan and Zuckerman's work "Randomness is Linear in Space" (*J. of Comp. and Sys. Sci.*, Vol. 52 (1), pages 43–52, 1996).

computation but rather to *state information* passed between two (computationally unbounded) phases. Thus, one may prefer to analyze this setting in terms of information theory. Either way, the computational aspect of this setting is extremely simple. In contrast, most work in Modern Cryptography refers to highly complex computational questions (e.g., the time-complexity of various computational tasks), which are only superficially understood at the present. This difference in the nature of the computational aspects in question yields a big difference in the technical aspects of the relevant works. But there is also a conceptual difference, to be discussed next.

Most work in Modern Cryptography refers to time-bounded adversaries, where the time-bound corresponds to "real time" due to the inherent limitations on the computing resources available to the adversary (at the present as well as in the relevant future). In other words, a successful break of a secure system requires so many computation steps that it is unlikely to occur in the relevant future (e.g., in the current century). However (typically), in the time-bounded model, a successful break will eventually occur (but, most probably, not in our life-time). In contrast (typically), in the bounded-storage model, if at the present the adversary has limited storage then security will be preserved throughout eternity.

On the importance of scale. Space-bounded adversaries were studied before (mainly in the context of zero-knowledge proofs).[3] A key difference between these prior works and the bounded-storage model is in *scale* (i.e., the relation of the space-bound to other parameters). Whereas these prior works have studied adversaries having a very small amount of storage (e.g., logarithmic in the security parameter), the bounded-storage model deals with adversaries having a huge amount of storage that need to deal with a (slightly) larger amount of information (or random noise). Indeed, a change in scale may seem as something very minor (or technical), but often it opens the door to totally different applications and implications. Such has been the case also in the domains of derandomization and PCP:

1. Pairwise-independent sample spaces allow to generate many samples at the (randomness-complexity) cost of generating two samples. When

[3] See Dwork and Stockmeyer's work "Zero-Knowledge With Finite State Verifiers" (in *Crypto88*, pages 71–75, 1988) and Kilian's work "Zero-knowledge with Log-Space Verifiers" (in *29th FOCS*, pages 25–35, 1988). In contrast, the bounded-storage model was introduced in Maurer's paper "Secret Key Agreement by Public Discussion from Common Information" (*IEEE Trans. on Inform. Th.* , Vol. 39 (3), pages 733–742, 1993).

the samples themselves are taken from a huge space, the impact of this discovery is a saving in the randomness-complexity. But when the samples themselves are taken from small sets, such pairwise-independent sample spaces yields a full derandomization (of the computation that relies on such a sequence of samples).[4]

2. Multi-prover interactive proofs (MIPs) were first shown to provide an alternative characterization of the class $\mathcal{NEXP}$. Their relevance to the more fundamental class $\mathcal{NP}$ was demonstrated by a scale-down, and has yield alternative characterizations of the class $\mathcal{NP}$. These characterizations have been the focus of all subsequent exciting developments regarding PCP and the intractability of approximation (of $\mathcal{NP}$-hard problems).[5]

Interestingly, whereas in these two examples, exciting developments followed by reducing the scale, in the case of the bounded-storage model the exciting developments followed by enlarging the scale.

[4] The work of Chor and Goldreich "On the Power of Two–Point Based Sampling" (*Jour. of Complexity*, Vol 5, 1989, pages 96–106) suggests to use such samples for approximating the average of a function defined over a huge set. Thus, the impact of their work was confined to saving (or "recycling") random bits. In contrast, Luby's (contemporary) work "A Simple Parallel Algorithm for the Maximal Independent Set Problem" (*SIAM J. on Comput.*, Vol. 15 (4), pages 1036–1053, November 1986) deals with samples taken from a relatively small set (e.g., the edge-set of a given graph), and yields a full derandomization of a simple randomized algorithm for a natural task. Indeed, Luby's work has inspired a vast amount of further research.

[5] The MIP characterization of $\mathcal{NEXP}$ was proven by Babai, Fortnow, and Lund in "Non-Deterministic Exponential Time has Two-Prover Interactive Protocols" (*Computational Complexity*, Vol. 1 (1), pages 3–40, 1991). Two alternative scale-downs of this result (to $\mathcal{NP}$) were shown in subsequent works: Babai, Fortnow, Levin, and Szegedy's "Checking Computations in Polylogarithmic Time" (in *23rd STOC*, pages 21–31, 1991) and Feige, Goldwasser, Lovász, Safra, and Szegedy's "Approximating Clique is almost NP-complete" (*J. of the ACM*, Vol. 43, pages 268–292, 1996).