

The Bare Bounded-Storage Model: The Tight Bound on the Storage Requirement for Key Agreement

Stefan Dziembowski and Ueli Maurer, *Fellow, IEEE*

Abstract

In the bounded-storage model (BSM) for information-theoretically secure encryption and key-agreement one makes use of a random string R whose length t is greater than the assumed bound s on the adversary Eve's storage capacity. The legitimate parties, Alice and Bob, execute a protocol, over an authenticated channel accessible to Eve, to generate a secret key K about which Eve has essentially no information even if she has infinite computing power. The string R is either assumed to be accessible to all parties or communicated publicly from Alice to Bob. While in the BSM one often assumes that Alice and Bob initially share a short secret key, and the goal of the protocol is to generate a much longer key, we consider in this paper the *bare* BSM without any initially shared secret key. It is proved that in the bare BSM, secret key agreement is impossible unless Alice and Bob have themselves very high storage capacity, namely $O(\sqrt{t})$. This proves the optimality of a scheme proposed by Cachin and Maurer.

I. INTRODUCTION

The bounded-storage model, proposed initially by Maurer in 1992 [14], is an approach to achieving provable security of cryptographic schemes even against an adversary with unlimited computational resources. This is called unconditional or information-theoretic security. The only assumption is that the adversary's storage capacity is bounded, say by s bits, where s can be very large. No computational hardness assumption is needed. The basic idea is to assume that a random t -bit string R is either temporarily available to the public (e.g. the signal of a deep space radio source) or broadcast by a satellite or by one of the legitimate parties. If $s < t$, then the adversary, called Eve, can store only partial information about R , but she is allowed to apply an arbitrary function $f : \{0, 1\}^t \rightarrow \{0, 1\}^s$ to R in order to compute the value she stores. No assumption about the feasibility of computing f is made.

The legitimate parties, called Alice and Bob, can each access a small fraction of the string R and execute a protocol, over an authenticated channel accessible to Eve, to generate a secret key K about which Eve has essentially no information, even if she has infinite computing power, and no matter which function f she applied. In the BSM one usually assumes that Alice and Bob initially share a short secret key that determines which bits of R they need to access and how they combine the accessed bits to result in the secret key K . In this model, which we call the standard BSM, the goal of a key-agreement protocol is that the derived key K is much longer than the initial key; in other words, the goal is key expansion rather than key generation. A long sequence of papers on key expansion [14], [2], [1], [9], [10], [13], [17] has led from partial security proofs (for special adversary strategies) to complete security proofs, and to the understanding that a scheme secure in the BSM is a special type of randomness extractor.

One can also consider a model, which we call the *bare* BSM, where Alice and Bob share no secret key initially. This model was first considered by Cachin and Maurer [5] who proposed a scheme in this model which requires Alice and Bob to each access $O(\sqrt{t})$ bits of R , much more than in the standard BSM with a short secret key. In this paper we prove that this is essentially optimal, i.e., that no secure key-agreement protocol for the bare BSM exists in which Alice and Bob access fewer than $O(\sqrt{t})$ bits of R . Such lower bound proofs, apart from being of general scientific interest, are important because they prevent the search for schemes that do not exist.

The BSM was also studied in the context of oblivious transfer [4], [7] and time-stamping [16].

II. THE BARE BOUNDED-STORAGE MODEL AND THE CACHIN-MAURER SCHEME

Key agreement in the BSM, from the adversary's viewpoint, consists of two phases.

In the first phase, the string R is available to all parties. Alice and Bob execute a protocol over a public channel, resulting in transcript T which Eve obtains. Then, based on the transcript, Alice and Bob each store some information about R . The protocol can be randomized, where R_A and R_B denote their respective (independent) random strings. More precisely, Alice stores $M_A = f_A(R, T, R_A)$, and Bob stores $M_B = f_B(R, T, R_B)$, for some functions f_A and f_B . Eve also stores some information $M_E = f_E(R, T, R_E)$ about R , where R_E denotes her randomness (which is, of course, independent of (R_A, R_B)).

In the second phase, R has disappeared. Alice and Bob execute a second (probabilistic) protocol based on the stored values M_A and M_B , resulting in a second transcript T' . Then Alice and Bob compute a secret key, K_A and K_B , respectively. It is not necessary to formalize this further, i.e., to make the functions used to compute K_A and K_B explicit.

The two security requirements are:

- (i) *Correctness*: The probability $P(K_A \neq K_B)$ that the keys are different should be negligible.
- (ii) *Secrecy*: The amount of information, $I(K_A; M_E T')$, obtained by Eve about the secret key (say K_A), must be negligible.

A scheme for key-agreement in the bare BSM was proposed by Cachin and Maurer in [5]. In their protocol, both Alice and Bob store an (independent) random subset of r bits of R , where r is on the order of \sqrt{t} . After R has disappeared for all parties, they publicly agree on which bits they have both stored. With very high probability, Eve has only partial information about these bits, and therefore Alice and Bob can apply privacy amplification (i.e., randomness extraction using a strong extractor with a public extractor parameter) to distill an essentially perfect key K . We prove in Section III that the protocol of [5] is essentially optimal.

III. LIMITATIONS OF KEY-AGREEMENT IN THE BARE BSM

A. Statement of the Lower Bound

We prove the following result, which shows that the practicality of such an approach without shared initial key is inherently limited: Alice or Bob must have storage capacity around \sqrt{s} . The proof is given in Section III-B. Let h be the binary entropy function defined as $h(p) = -p \log_2(p) - (1-p) \log_2(1-p)$.

Theorem 1: For any key-agreement protocol secure in the BSM for which $I(K_A; M_E T') \leq \delta$ and $P(K_A \neq K_B) \leq \epsilon$, the entropy of the secret key K_A generated by Alice is upper bounded by

$$H(K_A) \leq \frac{s_A s_B}{s} + h(\epsilon) + \epsilon s_A + \delta, \quad (1)$$

where s_A and s_B are Alice's and Bob's required storage capacities, respectively, and s is Eve's assumed storage capacity.

Observe that for small ϵ and δ , the right hand side of (1) becomes approximately equal to $(s_A s_B)/s$, and hence in any secure key agreement at least one of the parties needs to have memory of a size at least \sqrt{s} .

We note that this bound also implies a bound on the memory of the adversary in the protocol for the oblivious transfer in the bounded-storage model.¹ Namely, if the memory of the honest parties is s_A , then the memory of a cheating party has to be much smaller than s_A^2 . This shows that the protocol of [7] is essentially optimal and answers the question posed in [7], [8].

B. Proof of Theorem 1

Definition 1: A list Z_0, \dots, Z_n of random variables are *symmetric with respect to a random variable Y* if for every two sequences i_1, \dots, i_w and i'_1, \dots, i'_w of distinct indices we have

$$P_{Y, Z_{i_1}, \dots, Z_{i_w}}(y, z_1, \dots, z_w) = P_{Y, Z_{i'_1}, \dots, Z_{i'_w}}(y, z_1, \dots, z_w), \quad (2)$$

for all y, z_1, \dots, z_w .

¹This is because there exists a black-box reduction of the key-agreement problem to the oblivious transfer problem [12]. (It is easy to see that the reduction of [12] works in the bounded-storage model.)

In other words, the distribution of $(Y, Z_{i_1}, \dots, Z_{i_w})$ does not depend on the choice of the indices i_1, \dots, i_w .
Lemma 1: If Z_0, \dots, Z_n are symmetric with respect to Y , then there exists $i \in \{0, \dots, n\}$ such that

$$I(Y; Z_0 | Z_1 \cdots Z_i) \leq \frac{H(Y)}{n+1}.$$

Proof: The chain rule for conditional information² implies that

$$\sum_{i=0}^n I(Y; Z_i | Z_{i-1}, \dots, Z_0) = I(Y; Z_0, \dots, Z_n), \quad (3)$$

which is at most $H(Y)$. Therefore there must exist i such that

$$\frac{H(Y)}{n+1} \geq I(Y; Z_i | Z_{i-1}, \dots, Z_0)$$

By the symmetry condition (2) this last value can be replaced by $I(Y; Z_0 | Z_1, \dots, Z_i)$. This completes the proof. \blacksquare

A simple example of such symmetric variables is given below (we will use it later in the proof of the theorem).

Observation 1: Let Y and Z be random variables. Suppose the random variables Z_1, \dots, Z_n are sampled independently, each according to the distribution $P_{Z|Y}$. Then Z, Z_1, \dots, Z_n are symmetric with respect to Y .

The following observation will also be useful.

Observation 2: If Z_0, \dots, Z_n are symmetric with respect to Y then for an arbitrary function g the variables Z_0, \dots, Z_n are symmetric with respect to $g(Y)$.

Proof: For every y' from the domain of g , all sequences i_1, \dots, i_w and i'_1, \dots, i'_w of distinct indices, and every sequence z_1, \dots, z_w we have

$$\begin{aligned} P_{g(Y), Z_{i_1}, \dots, Z_{i_w}}(y', z_{i_1}, \dots, z_{i_w}) &= \\ \sum_{y: g(y)=y'} P_{Y, Z_{i_1}, \dots, Z_{i_w}}(y, z_{i_1}, \dots, z_{i_w}) &= \\ \sum_{y: g(y)=y'} P_{Y, Z_{i'_1}, \dots, Z_{i'_w}}(y, z_{i'_1}, \dots, z_{i'_w}) &= \\ P_{g(Y), Z_{i'_1}, \dots, Z_{i'_w}}(y', z_{i_1}, \dots, z_{i_w}), & \end{aligned} \quad (4)$$

where (4) follows from the assumption that Z_0, \dots, Z_n are symmetric with respect to Y . \blacksquare

To prove Theorem 1, recall that s_A, s_B , and s are the storage capacities of Alice, Bob, and Eve, respectively. We have to specify a strategy for Eve to store information (i.e., the function f_E). Such an admissible strategy is the following. For the fixed observed randomizer $R = r$ and transcript $T = t$, consider $\lfloor s/s_B \rfloor$ independent copies $M_B^1, \dots, M_B^{\lfloor s/s_B \rfloor}$ of what Bob stores, sampled independently according to the distribution $P_{M_B | R=r, T=t}$. (Clearly such sampling can be done by a computationally-unbounded Eve.)

Lemma 2: The random variables $M_B, M_B^1, \dots, M_B^{\lfloor s/s_B \rfloor}$ are symmetric with respect to M_A .

Proof: Recall that M_A is a randomized function of (R, T) , namely $M_A = f_A(R, T, R_A)$ for a random R_A . By Observation 1 variables $M_B, M_B^1, \dots, M_B^{\lfloor s/s_B \rfloor}$ are symmetric with respect to (R, T) , and hence also with respect to (R, T, R_A) since $R_A \rightarrow (R, T) \rightarrow M_B, M_B^1, \dots, M_B^{\lfloor s/s_B \rfloor}$ form a Markov chain. Thus by Observation 2 the random variables $M_B, M_B^1, \dots, M_B^{\lfloor s/s_B \rfloor}$ are symmetric also with respect to $M_A = f_A(R, T, R_A)$. \blacksquare

Hence Lemma 1 implies that there exists $i \in \{0, \dots, \lfloor s/s_B \rfloor\}$ such that

$$I(M_A; M_B | M_B^1, \dots, M_B^i) \leq \frac{H(M_A)}{\lfloor \frac{s}{s_B} \rfloor + 1} \leq \frac{H(M_A)}{\frac{s}{s_B}} \leq \frac{s_A s_B}{s}.$$

²Recall that the chain rule for information (see eg. [6], Theorem 2.5.2) states that for arbitrary random variables V_1, \dots, V_n , and U we have

$$I(U; V_0, \dots, V_n) = \sum_{i=0}^n I(U; V_i | V_{i-1}, \dots, V_0)$$

The last step follows from $H(M_A) \leq s_A$. Clearly an infinitely powerful Eve can compute such an index i . We hence assume that Eve stores $M_E := M_B^1, \dots, M_B^i$.³ Now we can apply Theorem 1 in [15] which considers exactly this setting, where Alice, Bob, and Eve have some random variables M_A , M_B , and M_E , respectively, jointly distributed according to some distribution $P_{M_A M_B M_E}$. The theorem states that the entropy of a secret key K that can be generated by public discussion is upper bounded as

$$H(K_A) \leq \underbrace{I(M_A; M_B | M_E)}_{\leq \frac{s_A s_B}{s}} + H(K_A | K_B) + \underbrace{I(K_A; M_E T^i)}_{\leq \delta}$$

Now, by Fano's lemma (cf. [3], p. 156)

$$H(K_A | K_B) \leq h(\epsilon) + \epsilon \log_2(2^{s_A} - 1) \leq h(\epsilon) + \epsilon s_A,$$

and we obtain (1). This concludes the proof of Theorem 1.

ACKNOWLEDGEMENTS

We would like to thank Louis Salvail and Christian Schaffner for pointing out an error in the proof stated in [11].

REFERENCES

- [1] Y. Aumann, Y. Z. Ding, and M. O. Rabin. Everlasting security in the bounded storage model. *IEEE Transactions on Information Theory*, 48(6):1668–1680, 2002.
- [2] Y. Aumann and M. O. Rabin. Information theoretically secure communication in the limited storage space model. In *Advances in Cryptology - CRYPTO '99*, pages 65–79, 1999.
- [3] R. E. Blahut. Principles and practice of information theory, 1987 Addison-Wesley Longman Publishing Co., Inc.
- [4] C. Cachin, C. Crepeau, and S. Marcil. Oblivious Transfer with a Memory Bounded Receiver. In Proc. of 39th FOCS, 1998.
- [5] C. Cachin and U. Maurer. Unconditional security against memory-bounded adversaries. In *Advances in Cryptology - CRYPTO '97*, pages 292–306, 1997.
- [6] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley and Sons, Inc., 1991.
- [7] Y. Z. Ding. Oblivious transfer in the bounded storage model. In *Advances in Cryptology - CRYPTO 2001*, pages 155–170, 2001.
- [8] Y. Z. Ding. *Provable Everlasting Security in the Bounded Storage Model*. PhD thesis, Harvard University, 2001.
- [9] Y. Z. Ding and M. O. Rabin. Hyper-encryption and everlasting security. In *STACS 2002, 19th Annual Symposium on Theoretical Aspects of Computer Science*, pages 1–26, 2002.
- [10] S. Dziembowski and U. Maurer. Tight security proofs for the bounded-storage model. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pages 341–350, 2002.
- [11] S. Dziembowski and U. Maurer. On generating the initial key in the bounded-storage model. In *Advances in Cryptology - EUROCRYPT 2004*, pages 126–137, 2004.
- [12] Y. Gertner, S. Kannan, T. Malkin, O. Reingold, and M. Viswanathan. Relationship between public key encryption and oblivious transfer. In *41st Annual Symposium on Foundations of Computer Science*, pages 325–339, 2000.
- [13] C. Lu. Hyper-encryption against space-bounded adversaries from on-line strong extractors. In *Advances in Cryptology - CRYPTO 2002*, pages 257–271, 2002.
- [14] U. Maurer. Conditionally-perfect secrecy and a provably-secure randomized cipher. *Journal of Cryptology*, 5(1):53–66, 1992.
- [15] U. Maurer. Secret key agreement by public discussion from common information. *IEEE Transactions on Information Theory*, 39(3):733–742, 1993.
- [16] T. Moran, R. Shaltiel, and A. Ta-Shma. Non-interactive Timestamping in the Bounded Storage Model. In *CRYPTO 2004*, volume 3152 of LNCS, pages 460–476. Springer, 2004.
- [17] S. Vadhan. On constructing locally computable extractors and cryptosystems in the bounded storage model. In *Advances in Cryptology - CRYPTO 2003*, pages 61–77, 2003.

³It may perhaps look a bit counter-intuitive that, for the sake of this proof, Eve does not necessarily store as many values as she could fit in her memory, i.e. set $M_E := M_B^{\lfloor s/s_B \rfloor}$. However, in principle it can be the case that $I(M_A; M_B | M_E^{\lfloor s/s_B \rfloor}) > I(M_A; M_B | M_E^i)$ (for $i < \lfloor s/s_B \rfloor$) because conditioning may actually increase a mutual information between random variables, i.e., $I(U; V) < I(U; V | W)$ is possible.