CSCI 1510 Capstone Project

Giovanni Pittalis
Banner ID: B01035414

May 15, 2018

Abstract

This is a report on Daniele Micciancio and Michael Walter's paper titled "On the Bit Security of Cryptogaphic Primitives" [1], which was presented at the 2018 Eurocrypt conference. This report explores the motivations of the paper, its main contributions, and how they affect the current state of cryptography research.

1 Introduction

The main problem that this paper tries to solve is the lack of a formal definition in cryptography literature when describing the security level of a cryptographic primitive by its number of "bits of security". When a certain cryptographic primitive is said to have n bits of security, the intuitive meaning is that it is as secure as a perfect cryptographic primitive with an n-bit key. In other words, if we define a generic attack cost T (which could be a combination of runtime, space usage, or other measure), any attacker would need to incur a cost of $T > 2^n$. Alternatively, an efficient attack would achieve a success probability ϵ of at most 2^{-n}. These two intuitive definitions can be combined to have that for any attack with cost T and success probability ϵ, it must be that $\frac{T}{\epsilon} > 2^n$. Even though this provides a qualitative understanding of bit security while also providing a concrete measure to compare different primitives, it is by no means a formal definition and it creates problems in certain situations, some of which are mentioned below:

- PRGs bit security. Given a PRG P with seed length n, the literature indicates that P cannot provide more than $\frac{n}{2}$ bits of security under the current definition. This is counter-intuitive since nothing prevents a n-bit seed from providing n bits of security, thus it should be possible for a PRG using such a seed to also provide n bits of security.

- Approximate samplers. When implementing a cryptographic scheme that requires sampling a distribution, using 53-bit precision floating point numbers leads to a 2^{-53} bound on statistical distance, which results in a 53-bit security on the scheme. Using the conventional definition of bit security, recent literature proved better bounds only for search primitives, thus it seems that decision primitives require higher precision than search primitives in order to obtain the same level of bit security. This is counter-intuitive since both search and decision primitives that rely on sampling distributions can leak information about the secret when the distribution is approximated up to a certain precision.

In addition, the definition does not generalize well from search primitives to decision primitives. The former are the primitives where an attacker tries to recover a secret such as a key, while in the latter the attacker tries to decide whether a secret bit is 0 or 1, such as indistinguishability problems. In search primitives, the number n of bits of security is understood as the logarithm of $\frac{T}{\epsilon}$ as described above, while in decision primitives the concept of bit security is not entirely clear. This is because the success probability is usually substituted by the distinguishing advantage, which is the probability that the attacker successfully guesses the correct hidden bit over the trivial probability $\frac{1}{2}$ of a random guess. Thus, the bit security for decision primitives is usually understood as the logarithm of $\frac{T}{\delta}$ where δ is the distinguishing advantage, which makes it different from the definition of bit security for search primitives.

2 Results

2.1 Advantage and Bit Security Definitions

The paper considers attackers that output a special "don't know" symbol $\perp$, which allows for a more comprehensive analysis of bit security. Given this, α is defined as the probability that the attacker doesn't output $\perp$ in a certain security game, and β is defined as the probability that, given the attacker did not output $\perp$, it correctly identifies the secret (for both search and decision primitives). Given this, the (conditional) distinguishing advantage of an attacker for decision primitives is given by $\delta = 2\beta - 1$.

In order to introduce the definitions of advantage and bit security, we to define the concept of security games as presented by the paper.

Definition 1. An n-bit *security game* is played by an adversary $\mathcal{A}$ interacting with a challenger X. At the beginning of the game, the challenger chooses a secret x, represented by the random variable $X \in \{0,1\}^n$, from some distribution $\mathcal{D}_X$. At the end of the game, $\mathcal{A}$ outputs some value, which is represented by the random variable A. The goal of the adversary is to output a value a such that $R(x,a)$, where R is some relation. $\mathcal{A}$ may output a special symbol $\perp$ such that $R(x,\perp)$ and $\overline{R}(x,\perp)$ are both false.

Given this definition, decision primitives have a 1-bit associated security game, while search primitives have a $O(\kappa)$ associated security game, where κ is the security parameter. We note that the definition allows any relation R to identify a successful attack, while for most application it is enough to consider the identity relation, i.e. the attacker correctly guesses the secret x.

We are now ready to provide the definition of advantage.

Definition 2. For any security game with corresponding random variable X and $A(X)$, the adversary $\mathcal{A}$'s *advantage* is

$$\mathrm{adv}^{\mathcal{A}} = \frac{I(X;Y)}{H(X)} = 1 - \frac{H(X|Y)}{H(X)}$$

where $I(X;Y)$ is the mutual information between X and Y, $H(\cdot)$ is the Shannon entropy[1], and $Y(X,A)$ is the random variable with marginal distributions $Y_{x,a} = \{Y \mid X = x, A = a\}$ defined as

1. $Y_{x,\perp} = \perp$, for all x.
2. $Y_{x,a} = x$, for all $(x,a) \in R$.
3. $Y_{x,a} = \{x' \leftarrow \mathcal{D}_X \mid x' \neq x\}$, for all $(x,a) \in \overline{R}$.

This definition is trying to measure the amount of information that the adversary learns about the secret. This could be achieved by considering the mutual information between the secret's random variable X and the adversary's output's random variable A. However, this wouldn't directly indicate when the adversary correctly guesses the secret, thus succeeding in the security game, as opposed to when it doesn't, since the adversary's output could reveal a lot of information about X without succeeding in the game. Hence the use of the random variable Y, which is clearly defined to handle the three possible cases with the adversary's output: either it is $\perp$, or it is a guess that is either correct or incorrect. Given this definition, it is then straightforward to define bit security, which follows from the intuition described in the introduction.

Definition 3. Let $T : \{\mathcal{A} \mid \mathcal{A} \text{ is any algorithm}\} \to \mathbb{Z}_+$ be a measure of resources that is linear under repetition, i.e. $T(k\mathcal{A}) = kT(\mathcal{A})$, where $k\mathcal{A}$ is the k time repetition of $\mathcal{A}$. For any primitive, we define its *bit security* as $\min_{\mathcal{A}} \log \frac{T(\mathcal{A})}{\mathrm{adv}^{\mathcal{A}}}$.

Even though Definition 2 captures very well the concept of advantage, it is cumbersome to use in practice. Therefore, the paper also provides a simpler definition in terms of the quantities we described in the introduction. For any adversary $\mathcal{A}$ playing a security game, let the *output probability* $\alpha^{\mathcal{A}} = \Pr[A \neq \perp]$ and the *conditional success probability* $\beta^{\mathcal{A}} = \Pr[R(X,A) \mid A \neq \perp]$. Then, in the context of decision primitives, let the

[1]Mutual information and Shannon entropy, including the conditional entropy of X given Y, $H(X|Y)$, are concepts from information theory. Their definitions can be found in [1], page 6.

conditional distinguishing advantage $\delta^{\mathcal{A}} = 2\beta^{\mathcal{A}} - 1$. Given these quantities, the paper contains a proof that the advantage in Definition 2 is equivalent to the following:

$$\mathrm{adv}^{\mathcal{A}} = \alpha \left(1 - \frac{(1-\beta)\log(2^n - 1) + H(\mathcal{B}_\beta)}{n}\right),$$

where $\mathcal{B}_\beta$ denotes the Bernoulli distribution with parameter β. This is still cumbersome, but it is much more useful if we try to approximate it. For large n, we have that $\mathrm{adv}^{\mathcal{A}}$ converges to $\alpha^{\mathcal{A}}\beta^{\mathcal{A}}$, which is exactly the success probability ϵ we would expect for search primitives. Plugging this in Definition 3, we obtain that the bit security is $\min_{\mathcal{A}} \frac{T(\mathcal{A})}{\alpha^{\mathcal{A}}\beta^{\mathcal{A}}}$, which is identical to the well-established definition of bit security for search primitives. On the other hand, for $n = 1$, we obtain by Taylor approximation that $\mathrm{adv}^{\mathcal{A}} \approx \alpha^{\mathcal{A}}(\delta^{\mathcal{A}})^2$. Plugging this in Definition 3, we obtain that the bit security is $\min_{\mathcal{A}} \frac{T(\mathcal{A})}{\alpha^{\mathcal{A}}(\delta^{\mathcal{A}})^2}$, which doesn't match the usual understanding of bit security for decision primitives that we described in the introduction. However, the paper notes that this matches an alternative definition put forward in a more specific context by Goldreich and Levin in previous literature. In addition, this definition directly solves the PRG bit security problem we described above, since the quadratic term (as opposed to a linear term) of the distinguishing advantage makes sure it is still possible to have a PRG with n bits of security given a n-bit seed. Therefore, the paper puts forward the following definition of advantage:

Definition 4. For a search game, the advantage of the adversary $\mathcal{A}$ is

$$\mathrm{adv}^{\mathcal{A}} = \alpha^{\mathcal{A}}\beta^{\mathcal{A}}$$

and for a decision game, it is

$$\mathrm{adv}^{\mathcal{A}} = \alpha^{\mathcal{A}}(\delta^{\mathcal{A}})^2$$

Although Definition 4 is not always equivalent to Definition 2, the paper contains a proof that it is equivalent in the context of bit security. In order to support this definition, the paper presents a series of technical results that are described in the following sections.

2.2 Security Reductions

The paper provides a series of tight reductions to prove a bound on the bit security of primitives that are built from other primitives. Since the new definitions for advantage and bit security diverge from the previous literature only on decision primitives, the reductions deal with constructions from search to decision, from decision to search, and from decision to decision primitives.

2.2.1 Search to Decision

The paper presents the Goldreich-Levin theorem as an example of this, which is a construction of a hardcore bit (a decision primitive) from a one-way function (a search

primitive). By using the same definition of bit security, Levin proved the following theorem:

Theorem 1. Let $f : \{0,1\}^n \to \{0,1\}^k$ be a κ-bit secure one-way function. Then $b(x,r) = \langle x, r\rangle \mod 2$ is a $(\kappa - O(\log n))$-bit secure hardcore bit for $g(x,r) = (f(x), r)$.

As expected, this provides evidence that, given a search primitive that is κ-bit secure, we can construct a decision primitive that is approximately also κ-bit secure. The paper notes that this proof is only possible by using a quadratic factor of $\delta^{\mathcal{A}}$ in the definition of advantage for decision primitives, as a linear factor would imply a dramatic loss of security in the reduction.

2.2.2 Decision to Search

The paper provides three examples of such constructions. The first is the well-known fact that PRGs (decision primitives) are also one-way functions (search primitives), as stated by the following theorem:

Theorem 2. If g is a PRG with κ-bit security, then it is also a $(\kappa - 4)$-bit secure one-way function.

This statement is possible to prove thanks to the fact that the adversary outputs $\perp$ in case of failure when inverting the PRG. In the standard reduction, the bit security of the resulting one-way function would incur a non-constant drop of $\log \alpha^{\mathcal{A}}$.

The second construction is a search primitive built from two indistinguishable distributions as follows:

Definition 5. Let $\mathcal{D}_0, \mathcal{D}_1$ be two distributions. We define the n-bit *secret recovery game* as the following n-bit security game: the challenger X chooses an n-bit secret $x \leftarrow \mathcal{U}(\{0,1\}^n)$ and sends the vector $\mathbf{c} = (c_i \leftarrow \mathcal{D}_{x_i})_{i \le n}$ to $\mathcal{A}$. The adversary $\mathcal{A}$ attempts to guess x.

In other words, the challenger chooses n samples from either $\mathcal{D}_0$ or $\mathcal{D}_1$, and the adversary has to guess the original distribution for each one of them in order to succeed. The paper presents a proof of the following theorem, which states that the given search primitive has approximately the same bit security as the original decision primitive:

Theorem 3. If the κ-bit secret recovery game is instantiated with two κ-bit secure indistinguishable distributions $\mathcal{D}_0$ and $\mathcal{D}_1$, and $\mathcal{D}_0$ is publicly sampleable, then it is $(\kappa - 1)$-bit secure.

The third construction is really a reduction which proves that IND-CCA secure encryption schemes are also message hiding with the same bit security, as stated by the following theorem:

Theorem 4. If a scheme with message space larger than 2^κ is κ-bit IND-CCA secure, it is k-bit message hiding.

We note that the proof for both of the last two examples are due to the fact that the adversary's advantage for decision primitives contains a quadratic factor of $\delta^{\mathcal{A}}$ as opposed to a linear one.

2.2.3 Decision to Decision

Given the new definition of advantage, the standard hybrid argument which is used to prove the security of reductions between two decision primitives needs to be revisited. In particular, the paper proves the following theorem:

Theorem 5. Let $\mathcal{H}_i$ be k distributions. If $\mathcal{H}_i$ and $\mathcal{H}_{i+1}$ are κ-bit indistinguishable for all i, then $\mathcal{H}_1$ and $\mathcal{H}_k$ are $(\kappa - 2(\log k + 1))$-bit indistinguishable.

The proof of this is due to the following lemma, whose proof is also provided in the paper:

Lemma 1. Let $\mathcal{H}_i$ be k distributions and $G_{i,j}$ be the indistinguishability game instantiated with $\mathcal{H}_i$ and $\mathcal{H}_j$. Further, let $\epsilon_{i,j} = \max_{\mathcal{A}} \mathrm{adv}^{\mathcal{A}}$ over all T-bounded adversaries $\mathcal{A}$ against $G_{i,j}$. Then $\epsilon_{1,k} \leq 3k \sum_{i=1}^{k-1} \epsilon_{i,i+1}$.

We note that this is very similar to the standard hybrid argument, except that the advantage of the adversary in distinguishing between $\mathcal{H}_1$ and $\mathcal{H}_k$ gains an additional factor of $3k$ in front of the sum of advantages. However, the paper argues that this is not a significant loss in bit security, since it only affects the constant in front of the $\log k$ term in Theorem 5, which would be still present given the standard hybrid argument.

As the final technical result, this paper tries to solve the problem with approximate samplers described in the introduction, by extending other results from the literature. Together with those results, the theorem proved by this paper implies that approximating a distribution with relative error bounded by $2^{-\kappa/2}$ allows to preserve almost all of κ bits of security. This means that by implementing distribution sampling by using limited precision 53-bit floating point numbers, it is still possible to obtain 100 bits or higher levels of security, whose proof had so far eluded the attempts of the research community.

References

[1] Daniele Micciancio and Michael Walter. On the Bit Security of Cryptographic Primitives. *IACR Cryptology ePrint Archive*, 2018.