Randomness in Cryptography February 14, 2013

Lecture 6: Encryption/Extraction

Lecturer: Yevgeniy Dodis *Scribe: Zahra Jafargholi*

In the previous lectures we showed that weak/block/SV/ BCS sources are not enough for traditional privacy, (not differential privacy). More historically people were interested in the study of imperfect sources; can we extract good randomness deterministically? No! Can we have privacy without deterministic extraction? In today's lecture we study a number of approaches to answer this question.

1 Setting

Let $K \in \{0,1\}^n$ be a key distribution and let $K \in S$, where S is a family of distributions. From now on we call S a *source*. We want to use K as a secret key for some encryption scheme $(\mathsf{Enc}, \mathsf{Dec})$ to encrypt message $m \in \{0,1\}^b$. We take two approaches.
Approach 1: Extract a traditional key R and use R to encrypt. $R = \mathsf{Ext}(K)$ where Ext is an extractor for distribution K. For example, if $|R| = b$ set $C = m \oplus R$, (One-Time pad). This approach is modular but can possibly be restrictive.
Approach 2: Encrypt directly using K and bypass extraction.
In particular, can we encrypt better without extraction?

Definition 1 Function $\mathsf{Ext} : \{0,1\}^n \to \{0,1\}^\ell$ is (S, ε)-*extractor* if for any distribution $K \in S$,

$$\mathsf{SD}(\mathsf{Ext}(K), U_\ell) \leq \varepsilon.$$

$\diamondsuit$

If S has an (S, ε)-*extractor*, we say S is (ℓ, ε)-*extractable*. For example, $\mathsf{SV}(\gamma)$ sources are $(1, \gamma)$-extractable.

Definition 2 An $(\mathsf{Enc}, \mathsf{Dec})$ scheme is (S, δ)-*secure* on $\{0,1\}^b$, if for any distribution $K \in S$ and for all $m \in \{0,1\}^b$,

$$\mathsf{SD}(\mathsf{Enc}(K, m), \mathsf{Enc}(K, U_b)) \leq \delta.$$

$\diamondsuit$

If S has $(\mathsf{Enc}, \mathsf{Dec})$ that is (S, δ)-*secure* we say S is (b, δ)-*encryptable.*

Lemma 1 *If S is (b, ε)-extractable then S is (b, ε)-encryptable.*

In general, if Ext is (b, ε)-extractor for S and $(\mathsf{Enc}, \mathsf{Dec})$ is (S, δ)-secure with key $R = U_b$ then $\mathsf{Enc}'(K, m) = \mathsf{Enc}(\mathsf{Ext}(K), m)$ is $(b, \varepsilon + \delta)$-secure. To prove Lemma 1 we only note that One-Time Pad is $(b, 0)$-secure so if S is (b, ε)-extractable then S is (b, ε)-encryptable.

2 Encryptability implies Extractability

In this section we prove the converse to Lemma 1.

Theorem 1 *For all $\varepsilon > 0$ and $b > \log n + 2\log 1/\varepsilon$ if source S is (b,δ)-encryptable then,*

(a) S is $(b - 2\log 1/\varepsilon, \varepsilon + \delta)$-extractable.

(b) If the encryption scheme is efficient ($Poly(n)$), S is efficiently $(b - 2\log 1/\varepsilon - \log n - O(1), \varepsilon + \delta)$-extractable.

Usually n is the security parameter and δ is negligible in n, in which case Theorem 1 implies that encrypting $b = \omega(\log n)$ bits implies extracting $b(1 - O(n))$ bits, i.e. One-Time pad is universal when you are encrypting a nontrivial number of bits $b = \omega(\log n)$.

Lemma 2 *Let source S be (b,δ)-encryptable by encryption Enc, where $b > \log n + 2\log 1/\varepsilon$ and $\varepsilon > 0$. Let $\mathcal{C} := \{\mathsf{Enc}(k, U_b) | k \in \{0,1\}^n\}$. If $\mathcal{C}$ is (ℓ,ε)-extractable then S is $(\ell, \varepsilon + \delta)$-extractable.*

Proof: Let Ext be a $(\mathcal{C},\varepsilon)$-extractor, define $\mathsf{Ext}'(k) := \mathsf{Ext}(\mathsf{Enc}(k,1))$.

$$\forall K \in S, \mathsf{SD}(\mathsf{Ext}'(K), U_\ell) = \mathsf{SD}(\mathsf{Ext}(\mathsf{Enc}(K,1)), U_\ell).$$

by triangle inequality,

$$\mathsf{SD}(\mathsf{Ext}(\mathsf{Enc}(K,1)), U_\ell) \leq \mathsf{SD}(\mathsf{Ext}(\mathsf{Enc}(K,1)), \mathsf{Ext}(\mathsf{Enc}(K,U_b))) + \mathsf{SD}(\mathsf{Ext}(\mathsf{Enc}(K,U_b)), U_\ell).$$

Since Ext is deterministic,

$$\mathsf{SD}(\mathsf{Ext}(\mathsf{Enc}(K,1)), U_\ell) \leq \mathsf{SD}(\mathsf{Enc}(K,1), \mathsf{Enc}(K,U_b)) + \mathsf{SD}(\mathsf{Ext}(\mathsf{Enc}(K,U_b)), U_\ell | K).$$

$\mathsf{SD}(\mathsf{Enc}(K,1), \mathsf{Enc}(K,U_b)) < \delta$ by security of encryption Enc and $\mathsf{SD}(\mathsf{Enc}(K,U_b)), U_\ell) < \max_k(\mathsf{SD}(\mathsf{Ext}(\mathsf{Enc}(k,U_b)), U_\ell) < \varepsilon$ by assumption. Hence,

$$\forall K \in S, \mathsf{SD}(\mathsf{Ext}'(K), U_\ell) \leq \delta + \varepsilon.$$

□

Note 1. If Ext and Enc are efficient so is Ext'.
Note 2. If Ext is efficient then the computational security of Enc implies that the ℓ extracted bits are pseudorandom.

Having Lemma 2, it suffices to show $\mathcal{C}$ is extractable. $\mathcal{C}$ has two properties a) every $C_k \in \mathcal{C}$ is a b-source, (i.e. $\mathbf{H}_\infty(C_k) = b$) and b) $|\mathcal{C}| = 2^n$ is "small", unlike S which we don't know anything about!

Lemma 3 *For all $\varepsilon > 0$ and any source $\mathcal{C}$ that consists of 2^n distributions C (over $\{0,1\}^s$) with $\mathbf{H}_\infty(C) \geq b$ where $b > \log n + 2\log(1/\varepsilon)$ then,*

(a) $\mathcal{C}$ is $(b - 2\log 1/\varepsilon, \varepsilon)$-extractable.

(b) $\mathcal{C}$ is efficiently $(Poly(n, s))$ $(b - 2\log 1/\varepsilon - \log n - o(1), \varepsilon)$-extractable.

Note that Lemma 2 and Lemma 3 imply Theorem 1.

Proof: Part a) We are going to show $\mathsf{Ext}(\mathcal{C})$ exists. For part a) we do not care about the efficiency of Ext. We pick Ext at random from $F := \{f | f : \{0,1\}^s \to \{0,1\}^\ell\}$ where $\ell = b - 2\log 1/\varepsilon$ and we prove

$$\Pr_{f \overset{r}{\leftarrow} F}\left[\exists C \in \mathcal{C}, \mathsf{SD}\left(f(C), U_\ell\right) > \varepsilon\right] < 1.$$

By a union bound it is enough to show that for any $C \in \mathcal{C}$,

$$\Pr_{f \overset{r}{\leftarrow} F}\left[\mathsf{SD}\left(f(C), U_\ell\right) > \varepsilon\right] < 1/2^n. \tag{1}$$

Statement (1) is an interesting statement on its own and we will return to it later. Here we state another lemma that will imply statement (1), and by proving this next lemma we conclude lemma 3 part a). For Lemma 3 part b) we will use the same method except that f is chosen from a family of efficient functions. □

Lemma 4 *Let $L = 2^\ell$, $B = 2^b$ and let C be a distribution over $\{0,1\}^s$. If $\mathbf{H}_\infty(C) \geq b \geq \log n + \log 1/\varepsilon$ then*

(a) $\Pr_{f \overset{r}{\leftarrow} F}\left[\mathsf{SD}\left(f(C), U_\ell\right) > \varepsilon\right] < 2^L \cdot e^{-\varepsilon^2 B}$.

(b) $\Pr_{f \overset{r}{\leftarrow} F}\left[\mathsf{RD}\left(f(C), U_\ell\right) > \varepsilon\right] < L \cdot e^{-2\varepsilon^2 B/L}$.

Part a) implies Lemma 3 part a) when $\ell = b - 2\log 1/\varepsilon$ and Lemma 3 part a) implies Theorem 1 part a).

Part b) requires $\ell \approx b - 2\log 1/\varepsilon - \log b$ for the probability to be less than 1 and $\ell \approx b - 2\log 1/\varepsilon - \log n$ to survive the union bound in the proof of Lemma 3 part b).

Proof: Without loss of generality let's assume C is uniform over B values. Define $f(C)$ to be the distribution of throwing B balls into L bins at random. Every bin has B/L balls in expectation. Let X_i = number of balls in bin i.

Part a)

$$\mathsf{SD}(f(C), U_\ell) = 1/2^B \sum_i |X_i - B/L| = \max_{T \subset [L]} 1/B \left(\sum_{i \in T}(X_i - B/L)\right).$$

To show part a) we take a union bound over all $T \in [L]$. Fix T, let Y_j be 1 if ball j is subset T and 0 otherwise. By Chernoff bound,

$$\Pr_f\left[\text{Surplus of } T \text{ is } \geq \varepsilon B\right] = \Pr_f\left[\sum_{j=1}^{B} Y_j \geq \varepsilon B + \mathbb{E}\left(\sum_{j=1}^{B} Y_j\right)\right] \leq e^{\varepsilon^2 B}$$

Therefore,

$$\mathsf{SD}(f(C), U_\ell) \leq 2^L \cdot e^{\varepsilon^2 B}$$

Part b)

$$\mathsf{RD}(F(C), U_\ell) \approx \max_i (\frac{X_i}{B/L} - 1)$$

To show part b) we use a union bound over all L bins and again apply Chernoff bound to X_i knowing $\mathbb{E}(X_i) = B/L$. □

We can improve lemma 4 by choosing Ext from a smaller family of functions $\mathbb{F}$ where functions in $\mathbb{F}$ have limited independence. For part a) we need L-independence and for part b) we need ℓ-independence to be able to afford the union bounds for each part.

Definition 3 $\mathbb{F} := \{f : \{0,1\}^s \to \{0,1\}^\ell\}$ is t-wise independent if for all distinct $C_1, C_2, ...C_t \in \{0,1\}^s$, $(f(C_1), f(C_2), ..., f(C_t) \equiv (U_1, U_2, ..., U_t)$ where $f \overset{r}{\leftarrow} \mathbb{F}$ and U_is are uniform. ♢

Fact 1. There exist t-wise independent families that every function in them satisfy $|f| = O(ts)$ and $TIME(f) = \tilde{O}(ts)$, (i.e. all the function in the family are efficient).
Fact 2. There is a Chernoff bound for t-wise independent variables,

$$Pr\left[|X - \mathbb{E}[X]| \geq \varepsilon \cdot \mathbb{E}[X]\right] \leq \left(\frac{t}{4\varepsilon \cdot \mathbb{E}[X]}\right)^{t/2}$$

Where X is the sum of all the t-wise independent variables.

Now, for part b) we pick $t = \ell = b - \log(1/\varepsilon) - \log b - O(1)$ and we get the bound we needed. To summarize, we divided the proof of Theorem 1 into two parts; first we reduced the existence of an extractor for a large source S to the existence of an extractor for a smaller source $\mathcal{C}$ and in the second part we showed that a random $n-$wise independent function, with high probability, is a good extractor for $\mathcal{C}$.

Question 1 *The extractor* Ext *(and hence* Ext'*) is efficient but its construction is non-uniform. Can we make it uniform? Or can we show that no "uniform" construction is possible? For example show that for any $Poly(n)$-size oracle circuit* $\mathsf{Ext}(\cdot)$ *there exists* $(\mathsf{Enc}, \mathsf{Dec})$ *and a distribution K such that* Enc *is $(K, 0)$-secure but* $\mathsf{SD}(\mathsf{Ext}^{(\mathsf{Enc})}(K), U_1) \geq 1/10$*. In other words for any extractor function f there exists an encryption* Enc *that fools it.*

3 Almost Perfect Resilient Functions

Lemma 3 has many interesting applications. It shows the existence of function f that is deterministic and a non-uniform extractor for a useful source $\mathcal{C}$. Now we give an example of application of Lemma 3.

Definition 4 A function $f : \{0,1\}^s \to \{0,1\}^\ell$ is $(b, \varepsilon) - APRF$ (Almost Perfect Resilient Function) if for all subsets $I \subset [s]$ of size b and for all $y \in \{0,1\}^{s-b}$,

$$\mathsf{SD}(f(C(I, y)), U_\ell) \leq \varepsilon.$$

Where $C(I, y)$ denotes an s-bit string that is random in the b positions corresponding to I and is fixed to y in the reset of $s - b$ positions. $\diamondsuit$

You can think of $C(I, y)$ as tampering with the input of function f. The attacker can choose $s - b$ positions of the string and fix them to a value y but the rest of the bits will remain random. We want to argue that the output of the APRF function is still statistically close to random. If a function is $(b, 0)$-APRF, it is called b-RF (Resilient Function). Note that for any $b \geq 1$ there exists a b-RF with $\ell = 1$, namely parity.

Now we define a family of distributions called Bit-fixing sources and then will use Lemma 3 on these sources to show the existence of Almost Perfect Resilient functions. $\mathcal{C}_b := \{C(I, y) | \forall I \subset [s], |I| = b, y \in \{0,1\}^{s-b}\}$ is called a Bit-fixing source and has two important properties:

1. $\forall\ I, y, \quad \mathbf{H}_\infty(C(I, y)) = b$.
2. $|\mathcal{C}_b| \leq \binom{s}{b}\ 2^{s-b} \leq 3^s$.

By Lemma 3 part a), if $b \geq 2 \log 1/\varepsilon + \log s + O(1)$ we know there exists a (b, ε)-APRF with $\ell = b - 2\log(1/\varepsilon) - O(1)$. By Lemma 3 part b) we know that there exists an efficient (b, ε)-APRF with $\ell = b - 2\log(1/\varepsilon) - \log s - O(1)$.

Fact 3. If $\ell > 2 \log s$ and f is ℓ-RF then $b > s/2$. In other words, we need an entropy rate of more than $1/2$ to extract $2 \log s$ bits perfectly. This lower bound on the entropy rate is essentially tight, due to a coding constriction for a Resilient function. Next, we will show this coding construction.

DEFINITION 5 code $E : \{0,1\}^\ell \to \{0,1\}^s$ is a linear (s, ℓ, d)-code, if there exists a generator matrix $G_{\ell \times s}$ such that $E(X) = X^t \cdot G$ and for all $X \neq 0$, the Hamming weight of $E(X) \geq d$ $\diamondsuit$

Theorem 2 *Let $f_E(C) = G \cdot C$. If E is a linear (s, ℓ, d)-code then f_E is b-RF, where $b = s - d + 1$.*

Proof: By Singleton bound we have $\ell \leq s - d + 1$ therefore $\ell \leq b$. Also by Plotkin bound we know that $d \geq s/2$ if $\ell \geq 2 \log s$. The point is that any linear combination of rows of G has weight more than or equal to $d \geq s - b + 1$. Hence, after removing $s - b$ columns of G, still any linear combination of rows is non-zero, i.e. still linearly independent which means that the remaining columns have full rank (ℓ). Let G_1 be the matrix that is the same as G in the columns of G that correspond to the $s - b$ fixed positions and has zero for the rest of the entries and let $G_2 = G - G_1 \mod 2$. Note G_2 has column rank and row rank ℓ.

$$G \cdot C = G_1 \cdot C + G_2 \cdot C.$$

$G_1 \cdot C$ is fixed and $G_2 \cdot C$ is random ($\equiv U_\ell$).

□

Kaoru Kurosawa et al. showed how to beat this coding construction for (b, ε)-APRF, meaning how to extract more bits than the above construction but we still need $b \geq s/2$. One open problem is whether it is possible to construct an APRF with $b < s/2$ and $\ell \geq 2 \log s$.

4 ERF and AONT

So far we have considered the case where the attacker could tamper with the input of function f. Now we go from tampering to leakage. We mention two type of functions, *ERF* and *AONT*. First one is the Exposure Resilient Functions. Here the attacker A cannot tamper with the input of function f and the input C is random in $\{0,1\}^s$. What the attacker can do is to adaptively choose a bit position of C and learn the value of that bit and the attacker can do this for at most $s-b$ bits of C. We call $f:\{0,1\}^\ell \rightarrow \{0,1\}^s$ a (b,ε)-Exposure Resilient Function if

$$\mathsf{SD}(f(C), U_\ell | view(A)) \leq \varepsilon.$$

Lemma 5 *If f is (b,ε)-APRF then f is (b,ε)-ERF.*

AONT stands for All-Or-Nothing Transforms. Here we have a secret X of length ℓ and we want to store it in Z of size s bits in a way that if attacker A learns $s-b$ bits of Z, X is still secure. We call a function $f:\{0,1\}^\ell \rightarrow \{0,1\}^s$ to be (b,ε)-AONT if

$$\forall x_0, x_1 \in \{0,1\}^\ell,\ \mathsf{SD}(A(f(x_0), A(f(x_1)) \leq \varepsilon.$$

Lemma 6 *If f is (b,ε)-ERF then f is (b,ε)-AONT.*

References

[1] Yevgeniy Dodis, Adam Sahai, Adam Smith. On Perfect and Adaptive Security in Exposure-Resilient Cryptography . In *EUROCRYPT 2001.*

[2] Jesse Kamo, David Zuckerman. Deterministic Extractors for Bit-Fixig Sources and Exposure-Resilient Cryptography. In *SICOMP 2006*