Applied Cryptography

Symmetric Cryptography, Assignment 2, Wednesday, March 2, 2022

Remarks:

- Hand in your answers through Brightspace.
- Hand in format: PDF. Either hand-written and scanned in PDF, or typeset and converted to PDF. Please, **do not** submit photos, Word files, LaTeX source files, or similar.
- Assure that the name of **each** group member is **in** the document (not just in the file name).

Deadline: Wednesday, March 16, 23.59

Goals: After completing these exercises you should have understanding in arguing security of hash functions, and in the usage of sponge and duplex functions.

1. **(10 points)** In lecture 4 slides 11 (and also at Introduction to Cryptography), we have learned that the Merkle-Damgård hash function is collision resistant if the underlying compression function F is collision resistant. However, for this to hold, we *require* the initial state value to be a constant IV or the message length $\text{len}(M)$ to be encoded. Suppose that we would not do so, i.e., fill the initial state with message data and omitting the length encoding:

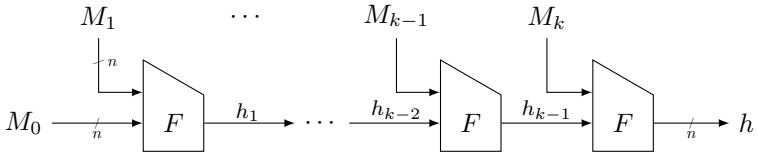

   Mount a generic collision attack against this mode.

2. **(10 points)** Merkle-Damgård is a *sequential* hashing mode, an alternative is tree-based hashing. This exercise is about proving collision resistance of a *simplified* tree-based hashing mode. Let $F : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n$ be a compression function, and consider the following hash function $\mathcal{H} : \{0,1\}^{4n} \to \{0,1\}^n$:

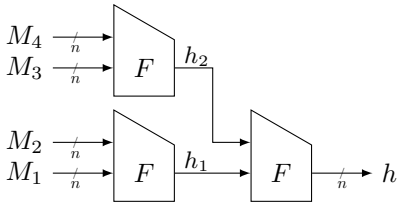

   Argue that $\mathcal{H}$ is collision resistant if F is collision resistant. Note that in the lecture we did not state a formal security model (for technical reasons); therefore, you do not have to derive a formal security bound in this exercise.

3. **(0 points, as this question already appeared in Introduction to Cryptography)** The sponge construction is proven to be indifferentiable from random with bound around $N^2/2^{c+1}$. The bound is, in fact, tight: it is possible to "break" the construction by making $N \approx 2^{c/2}$ evaluations of the permutation. Describe an attack that obtains a collision for the sponge construction with approximately $N \approx 2^{c/2}$ evaluations of the permutation. For simplicity, restrict your focus to data that is of length 3 blocks, where padding appends an encoding of the length into the last block:

$$\mathsf{pad}(M) = M \| 0^* \| \langle \mathsf{length}(M) \rangle_r \,.$$

This means that you can exactly consider the sponge construction as depicted on lecture 4 slide 25.

Remarks: the precise number of permutation calls in your attack may differ by a constant factor (this is not a problem), and you do not need to compute the exact success probability (intuition suffices).

4. **(10 points)** The sponge construction is proven to be indifferentiable from random with bound around $N^2/2^{c+1}$. The bound is, in fact, tight: it is possible to "break" the construction by making $N \approx 2^{c/2}$ evaluations of the permutation. However, it is not immediate to use such "break" to obtain second preimages for the construction.

   (a) Consider a sponge construction with $b = 320$, $c = 256$, $r = 64$, and $n = 256$. This particularly means that digests are generated through 4 squeezes. What is the average number of primitive evaluations N you need to make in order to find a second preimage? Explain your answer.

   (b) For the setting of question (a), suppose you are given a message M and its image $h \in \{0,1\}^n$. Describe informally an attack that obtains a second preimage for h with approximately N evaluations of the permutation, where N is your answer to question (a).

   (c) Consider a sponge construction with $b = 300$, $c = 256$, $r = 44$, and $n = 88$. This particularly means that digests are generated through 2 squeezes. What is the average number of primitive evaluations N you need to make in order to find a second preimage? Explain your answer.

   (d) For the setting of question (c), suppose you are given a message M and its image $h \in \{0,1\}^n$. Describe informally an attack that obtains a second preimage for h with approximately N evaluations of the permutation, where N is your answer to question (c).

5. **(10 points)** Consider SpongeWrap from lecture 5 slide 10. This construction is only a secure authenticated encryption scheme if the distinguisher cannot repeat nonces for encryption queries (it may repeat nonces for decryption queries, though).

   (a) Suppose above condition is violated and the distinguisher can repeat nonces for encryption queries. Describe an attacker that breaks the confidentiality of SpongeWrap in a constant number of queries.

   (b) Suppose we *omit* the domain separating bits 1/0 from SpongeWrap, and associated data and message are both padded with simple 10^*-padding (i.e., the last, incomplete, block of both A and M is padded with a 1 and a sufficient number of 0s to get an r-bit block). Describe an attacker that breaks the authenticity of SpongeWrap in a constant number of queries. (Note: the distinguisher is *not* allowed to repeat nonces for encryption queries.)

6. **(10 points)** This question is about the keyed sponge/duplex in a leaky setting. In this setting, we assume that every evaluation of P on secret state leaks λ bits of information about this state. This happens in a non-adaptive setting. I.e., if S is a secret state, and the attacker can eavesdrop an evaluation $P(S \oplus (M\|0^c))$, λ bits of the input $S \oplus (M\|0^c)$ leak to the adversary. If it can also eavesdrop an evaluation $P(S \oplus (M'\|0^c))$, for a different M', λ different bits (in the worst case) of the input $S \oplus (M'\|0^c)$ are leaked. If a permutation evaluation is repeated, this yields *no new* information. (The adversary would typically be assumed to have set up its eavesdropping attack so as to maximize the leakage, so the smartest possible attacker would, in above example, have learned 2λ bits of S.)

   Consider SpongeWrap from lecture 5 slide 10. For the sake of simplicity, we assume that $b = 3k$, $c = 2k$, and thus, $r = k$. K is a k-bit key, N a k-bit nonce, and $\mathsf{init}(K, N) = K\|N\|0^k$.

(a) Describe informally an attack that obtains the secret k-bit key, assuming λ bits leakage per primitive evaluation. How many construction queries does your attack make (in terms of k and λ)?

We reconsider SpongeWrap by *not* gluing together key and nonce, but *instead* inserting the nonce as A_1 (and move up associated data by one permutation call).

(b) Suppose the adversary can make 2 authenticated encryption evaluations for different nonces. How many bits of leakage does it learn of the key K? (Hint: it is important to remember that if a permutation is evaluated twice for the same input value, it does *not* leak 2λ bits but rather only λ bits.)

(c) It is nevertheless possible to recover the key by making $\lceil b/\lambda \rceil$ authenticated encryption evaluations. Describe informally an attack that obtains the secret key K with approximately this many number of construction evaluations.

(d) Consider the following stream encryption mode based on the sponge. As before, we have a k-bit key K and k-bit nonce N. This time, the nonce is cut into k bits $N_1 \| \cdots \| N_k$, that are processed one by one. Keystream generation happens with r-bit blocks at a time.

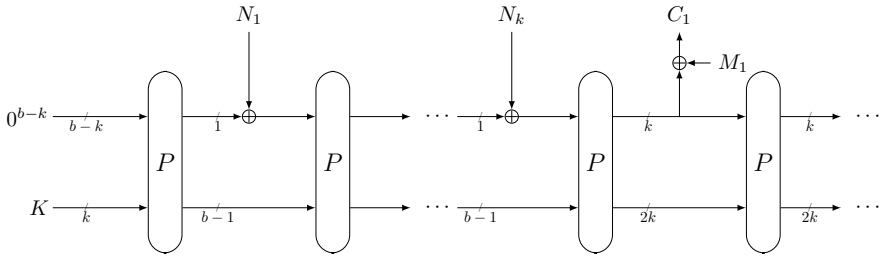

Explain informally why leakage is no problem here (neither at the absorbing phase nor at the squeezing phase).