A Note on Adversarial Online Complexity in Security Proofs of Duplex-Based Authenticated Encryption Modes

Charlotte Lefevre

Digital Security Group, Radboud University, Nijmegen, The Netherlands
charlotte.lefevre@ru.nl

Abstract. This note examines a nuance in the methods employed for counting the adversarial online complexity in the security proofs of duplex-based modes, with a focus on authenticated encryption. A recent study by Gilbert et al., reveals an attack on a broad class of duplex-based authenticated encryption modes. In particular, their approach to quantifying the adversarial online complexity, which capture realistic attack scenarios, includes certain queries in the count which are not in the security proofs. This note analyzes these differences and concludes that the attack of Gilbert et al, for certain parameter choices, matches the security bound.

1 Introduction

Permutation-based cryptography has gained significant popularity in recent years, with the sponge [BDPV07] and later the duplex [BDPA11] constructions introduced by Bertoni et al. The duplex construction, in particular, can be used to build authenticated encryption (AE) modes, such as `SpongeWrap` [BDPA11], `MonkeySpongeWrap` [Men23], `Cyclist` [DHP+20], and `Ascon` [DEMS21]. While the duplex construction offers a large range of possible applications, its inherent complexity often leads to complicated security bounds. Mennink [Men23] conducted a detailed analysis of these bounds, illustrating their versatility in diverse applications, including authenticated encryption. These complex bounds are occasionally prone to misinterpretation. For instance, the designers of `Xoodyak` [DHP+20] misinterpreted the security bound, leading to a security claim that was later broken by Gilbert et al. [GBKR23]

This note aims to discuss one specific aspect of duplex-based security proofs: adversarial online complexity. We will focus on the security bounds applied in the context of authenticated encryption particularly examining the methodology used by Gilbert et al. [GBKR23] as one example of counting adversarial online complexity. By analyzing their approach and its implications, we aim to understand how it aligns or diverges from the view of the security proofs.

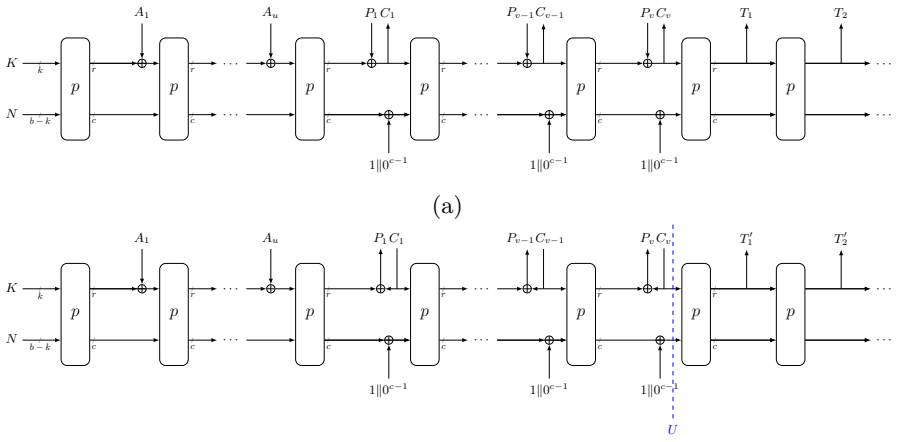

Fig. 1: Encryption (a) and Decryption (b) of `MonkeySpongeWrap` based on the permutation p. The associated data, plaintext and ciphertext are padded into r-bit blocks as $A_1 \parallel \cdots \parallel A_u \leftarrow \text{pad}(A)$, $P_1 \parallel \cdots \parallel P_v \leftarrow \text{pad}(P)$, $C_1 \parallel \cdots \parallel C_v \leftarrow \text{pad}(C)$. For encryption, the ciphertext is obtained by truncating $C_1 \parallel \cdots \parallel C_v$ to its first $|P|$ bits. For decryption, the plaintext is obtained by truncating $P_1 \parallel \cdots \parallel P_v$ to its first $|C|$ bits. The plaintext is returned if and only if the tag computed coincides with the tag given as input.

2 Preliminaries

2.1 Notation

This paragraph introduces useful notation. Let $b \in \mathbb{N}$. $\{0,1\}^b$ denotes the set of binary strings of size b bits, and $\{0,1\}^* = \bigcup_{a \in \mathbb{N}} \{0,1\}^a$. Given $X \in \{0,1\}^b$ and $n \in \mathbb{N}$ such that $n \leq b$, $\lfloor X \rfloor_n$ represents the rightmost n bits of X, and $\lceil X \rceil_n$ the leftmost n bits of X. $x \xleftarrow{\$} S$ denotes that x is sampled uniformly at random from a finite set S. $\perp$ denotes the bottom function, that given any tuple of inputs, returns the special symbol $\perp$. Given $m \in \mathbb{N}$, $\text{Perm}(m)$ denotes the set of permutations over $\{0,1\}^m$.

2.2 Duplex-Based Authenticated Encryption Modes

This note does not aim to provide a comprehensive overview of duplex-based AE modes. Instead, our focus is on describing one typical duplex-based AE mode that reflects practical design considerations and covers a broad range of use cases. Readers interested in an overview of the duplex construction are encouraged to explore the detailed work by Mennink [Men23]. We will describe the `MonkeySpongeWrap` authenticated encryption mode [Men23]. It is important to note, however, that the mode of `Ascon` cannot be described in terms

of `MonkeySpongeWrap` due to its unique key-blinding technique (see Section 9.3 of [Men23] for other examples).

`MonkeySpongeWrap` operates on a state of size b bits split as $b = r + c$. The r leftmost bits of the state will be referred to as the outerpart, and the rightmost c bits as the inner part. Let $\mathrm{pad}(M)$ be the function that appends to M a 1, and as many zeros as necessary to obtain a message of length multiple of r, and let $p \in \mathrm{Perm}(b)$ be a cryptographic permutation. The encryption function, instantiated with p, and based on the key $K \in \{0,1\}^k$ is denoted by $\mathcal{E}_K^p$. It takes as input a tuple $(N, P, A) \in \{0,1\}^n \times \{0,1\}^* \times \{0,1\}^*$, where N is the nonce, P the plaintext, and A the associated data. It outputs a ciphertext $C \in \{0,1\}^{|P|}$ and a tag $T \in \{0,1\}^t$. Conversely, the decryption function $\mathcal{D}_K^p$ takes as input $(N, C, A, T) \in \{0,1\}^n \times \{0,1\}^* \times \{0,1\}^t$. It outputs either the special symbol $\perp$, or a plaintext $P \in \{0,1\}^{|C|}$. The encryption and decryption functions are illustrated in Fig. 1a and Fig. 1b respectively.

2.3 Security Model

In the following the AE security of `MonkeySpongeWrap` is defined in the single-user setting. Let \$ denote a random function that, given a fresh input $P \in \{0,1\}^*$, outputs a random string of size $|P| + t$. The distinguisher $\mathcal{A}$ has oracle access to either $\left((\mathcal{E}_K^p, \mathcal{D}_K^p), p^{\pm}\right)$ in the real world, or to $\left((\$, \perp), p^{\pm}\right)$ in the ideal world. The security of `MonkeySpongeWrap` against $\mathcal{A}$ is given by

$$\mathrm{Adv}_{\mathtt{MSW}}(\mathcal{A}) = \Big|\mathbf{Pr}\left(K \xleftarrow{\$} \{0,1\}^k, p \xleftarrow{\$} \mathrm{Perm}(b) : \mathcal{A}^{\left((\mathcal{E}_K^p, \mathcal{D}_K^p), p^{\pm}\right)} \to 1\right) - \mathbf{Pr}\left(p \xleftarrow{\$} \mathrm{Perm}(b) : \mathcal{A}^{\left((\$, \perp), p^{\pm}\right)} \to 1\right)\Big|.$$

The adversary is restricted by two constraints. First, it is not allowed to make a decryption query with (N, C, A, T) such that (C, T) is the output of a previous encryption query of form (N, P, A). Then, the adversary is nonce-respecting, meaning that it never makes two different encryption queries with the same nonce. The adversary can be characterized by the following resources:

- q_P: number of permutation queries;
- q_E: number of encryption queries;
- q_D: number of decryption queries;
- σ_E: total block length of encryption queries;
- σ_D: total block length of decryption queries.

With these quantities, let $\mathrm{Adv}_{\mathtt{MSW}}(q_P, \sigma_D, \sigma_E, q_D, q_E)$ denote the maximum of $\mathrm{Adv}_{\mathtt{MSW}}(\mathcal{A})$, among all adversaries $\mathcal{A}$ with these resources.

The quantities σ_E and σ_D count the number of permutation queries induced by encryption (resp., decryption) queries needed in the world $\left((\mathcal{E}_K, \mathcal{D}_K), p^{\pm}\right)$. This note emphasizes a crucial point: *the quantities σ_E and σ_D do not double-count permutation evaluations due to repeated paths.* For instance, if the adversary makes two encryption queries with empty associated data and plaintexts

(after padding) $(P_1 \parallel P_2)$ and $(P_1 \parallel P_2')$, where $P_1, P_2, P_2' \in \{0,1\}^r$ and $P_2 \neq P_2'$, then the value of σ_E equals 3. However, if the permutation calls inside construction calls repeat in $((\mathcal{E}_K^p, \mathcal{D}_K^p), p^{\pm})$ due to an unlucky collision, then these are doubly counted (but considered as a bad event anyway). From a security proof perspective, this counting method is explained by the fact that one large part of the analysis consists of bounding the probability of certain bad events, which are triggered by fresh permutation calls. Therefore, doubly counting a repeating path, given the current approaches, would only result in a more lossy bound.

It is worth noting that this phenomena also applies to the more general bounds of the duplex. In that case, the online complexity is the number of *distinct* duplexing calls [DMA17, Men23, DM19], and doubly repeating paths are not counted twice either.

3 `MonkeySpongeWrap` Security and Gilbert et al. Attack

3.1 Simplified Security Bound of AE-based Duplex Modes

Soon after its introduction, the sponge has been proven to be indifferentiable [MRH04, CDMP05] from a random oracle [BDPA08]. In a bit more detail, when assuming that the underlying permutation p is random, no adversary making less than $2^{c/2}$ queries to p can differentiate the sponge from a random oracle with a non-negligible probability. This is a powerful result, and has been used among others to prove the security of `SpongeWrap` [BDPA11]. Subsequent research [MRV15, JLM14, JLM+19] showed that it is possible to achieve security beyond $c/2$ bits in keyed applications, and further optimizations of the keyed duplex can be made. This phenomena is reflected in the general bound of the duplex of Daemen et al [DMA17]. Mennink [Men23] used this bound to map the resources of the adversaries for the use-case of `MonkeySpongeWrap`. Let $\sigma = \sigma_D + \sigma_E$ and $q = q_E + q_D$. The simplified derived bound has form (in the single-user setting):

$$\mathrm{Adv}_{\mathtt{MSW}}(q_P, \sigma_D, \sigma_E, q_D, q_E) = \mathcal{O}\left(\frac{(q_P + q + \sigma)^2}{2^b} + \frac{\nu_{r,c}^{\sigma}(q_P + \sigma_D + q_E)}{2^c} + \frac{q^2}{2^{\min(b,c+k)}} + \frac{q_P}{2^k} + \frac{q_D}{2^t} + \frac{\binom{\sigma_D}{2}}{2^c} + \frac{\sigma_D q_P}{2^c} \right),$$

where $\nu_{r,c}^{\sigma}$ denotes a multi-collision limit function, oftenly small (see Section 4.2 of [Men23] for a detailed discussion). Similar bounds have been derived for specific modes, such as `NORX` [JLM14, JLM+19]. Note, the aforementioned expression is a simplification, in particular by grouping quadratic factors appearing in the term in 2^b (making this simplification lossy). However, in this discussion we will only focus on the term in $\frac{\sigma_D q_P}{2^c}$, which appears as such in the bound.

Because of the terms $\frac{\binom{\sigma_D}{2}}{2^c}$ and $\frac{\sigma_D q_P}{2^c}$, duplex-based AE modes cannot be strictly categorized as beyond the birthday bound in c[1]. The reason why these terms are quadratic in the capacity is that with decryption queries, the adversary can fix the outerpart of the state to a value of their choosing, thus making collisions on the inner part potentially dangerous. This property is exploited by Gilbert et al. [GBKR23] to mount their attack.

3.2 Attack of Gilbert et al.

The forgery attack from Gilbert et al [GBKR23] covers a broad class of duplex-based AE modes where the key is incorporated to the state once, during the initialization phase. This includes, among others, `SpongeWrap` and `MonkeySpongeWrap`. This attack needs only decryption and permutation queries, and targets the term $\frac{\sigma_D q_P}{2^c}$ in the security bound. Their attack exploit the fact that for a fixed $\beta \in \{0,1\}^r$, and a random permutation p, the function $F_\beta(x) = \lfloor p(\beta \parallel x) \rfloor_c$ behaves like a random function. A simplified overview of the attack is as follows (refer to [GBKR23] for a detailed description):

1. Offline phase: find a $\beta \in \{0,1\}^r$ such that the function F_β has a large component, but all paths in this component terminate in a small cycle. Then each element in the small cycle is seen as a candidate state value before computing the tag (called "U" in Fig. 1b). For each $x \in \{0,1\}^c$ in the cycle, and for $i = 1, \ldots, \lceil \frac{t}{r} \rceil$, compute $p^i(\beta \parallel x)$ to obtain a candidate tag T_x;
2. Online phase: fix a nonce N, and let $C = \beta^\ell$. All subsequent decryption queries will have the nonce N, ciphertext C, and empty associated data. In particular, the path induced by the permutation evaluations, and hence the tag to guess, is the same for all decryption queries. If ℓ is sufficiently large, with high probability the $\ell - 1$ successive permutation evaluations made to absorb C end up in the small cycle computed beforehand. In that case, submitting all of the decryption queries (N, C, ϵ, T_x) with all of the T_x tags computed beforehand allow the adversary to succeed a forgery.

In the first step, they propose to spend $\approx 2^{3c/4}$ permutation queries, which would give a cycle of length $\approx 2^{c/4}$ with high probability. Then, with $\ell \approx 2^{c/2}$ in the second phase, the attacks succeeds with a high probability. From this information, we can already deduce that $q_P \approx 2^{3c/4}$, $q_D \approx 2^{c/4}$.

Regarding the online complexity, they indicate a total online complexity of $2^{3c/4}$. The rationale is as follows: there are in total $q_D = 2^{c/4}$ decryption queries, each query requiring $2^{c/2}$ permutation calls, thus leading a total online complexity of $2^{3c/4}$. Let us call this quantity the attack-wise cost, and denote it by $\tilde{\sigma_D}$. $\tilde{\sigma_D}$ is a valid upperbound of the σ_D used in the security proofs, but could be decreased if we take the metrics used by the security proofs. As explained in Section 2.3, in the security proofs of the duplex construction, the blocks for encryption/decryption queries that repeat are not doubly counted. Therefore, in

[1] Arguably, one could imagine that in real-life scenarios, the system blocks if too much failed decryption queries are made.

the setting of this attack, only the first decryption query increments σ_D, as the subsequent decryption queries only change the tag. With the metrics of the security proof, $\sigma_D = 2^{c/2}$, and $q_D = 2^{c/4}$, which gives $q_P\sigma_D = 2^{5c/4}$. In particular, this parametrization does not match the security bound.

If we parametrize this attack differently, i.e., by being satisfied with a cycle in the first phase of size $2^{c/2}$, then the attack succeeds as well with high probability when $q_P \approx 2^{c/2}$, and presents a cost of $q_D \approx 2^{c/2}$, $\sigma_D \approx 2^{c/2}$, and $\tilde{\sigma_D} \approx 2^c$. While the attack-wise cost $\tilde{\sigma_D}$ is much higher in this parametrization, the proof-wise cost σ_D is lower. Notably, when the tag size and key length are larger than $c/2$, this parametrization of their attack matches the security bound. This observation does not enhance the attack per se but offers alternative compromises when considering the counting method used in security proofs.

To conclude, the attack from Gilbert et al., [GBKR23], for certain parameter sets, matches the security bounds of typical duplex-based AE modes. The gap here rather lies in the way queries are counted, as the security proofs assume that a permutation evaluation done with one path is counted only once. This assumption is unrealistic in light of real-life scenarios, as in practise past permutation evaluations are forgotten and re-made, and the current existing proofs do not account for that.

Acknowledgements

Thanks to the authors of [GBKR23] for the insightful discussions. Thanks as well to Bart Mennink for fruitful discussions, as well for his contribution to the tikzpicture. Charlotte Lefevre is supported by the Netherlands Organisation for Scientific Research (NWO) under grant OCENW.KLEIN.435.

References

BDPA08. Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. On the indifferentiability of the sponge construction. In Nigel P. Smart, editor, *Advances in Cryptology - EUROCRYPT 2008, 27th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Istanbul, Turkey, April 13-17, 2008. Proceedings*, volume 4965 of *Lecture Notes in Computer Science*, pages 181–197. Springer, 2008.

BDPA11. Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. Duplexing the sponge: Single-pass authenticated encryption and other applications. In Ali Miri and Serge Vaudenay, editors, *Selected Areas in Cryptography - 18th International Workshop, SAC 2011, Toronto, ON, Canada, August 11-12, 2011, Revised Selected Papers*, volume 7118 of *Lecture Notes in Computer Science*, pages 320–337. Springer, 2011.

BDPV07. Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. Sponge functions. Ecrypt Hash Workshop 2007, May 2007.

CDMP05. Jean-Sébastien Coron, Yevgeniy Dodis, Cécile Malinaud, and Prashant Puniya. Merkle-Damgård Revisited: How to Construct a Hash Function. In

Victor Shoup, editor, *Advances in Cryptology - CRYPTO 2005: 25th Annual International Cryptology Conference, Santa Barbara, California, USA, August 14-18, 2005, Proceedings*, volume 3621 of *Lecture Notes in Computer Science*, pages 430–448. Springer, 2005.

DEMS21. Christoph Dobraunig, Maria Eichlseder, Florian Mendel, and Martin Schläffer. Ascon v1.2. Winning Submission to NIST Lightweight Cryptography, 2021.

DHP+20. Joan Daemen, Seth Hoffert, Michaël Peeters, Gilles Van Assche, and Ronny Van Keer. Xoodyak, a lightweight cryptographic scheme. *IACR Trans. Symmetric Cryptol.*, 2020(S1):60–87, 2020.

DM19. Christoph Dobraunig and Bart Mennink. Leakage resilience of the duplex construction. In Steven D. Galbraith and Shiho Moriai, editors, *Advances in Cryptology - ASIACRYPT 2019 - 25th International Conference on the Theory and Application of Cryptology and Information Security, Kobe, Japan, December 8-12, 2019, Proceedings, Part III*, volume 11923 of *Lecture Notes in Computer Science*, pages 225–255. Springer, 2019.

DMA17. Joan Daemen, Bart Mennink, and Gilles Van Assche. Full-state keyed duplex with built-in multi-user support. In Tsuyoshi Takagi and Thomas Peyrin, editors, *Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part II*, volume 10625 of *Lecture Notes in Computer Science*, pages 606–637. Springer, 2017.

GBKR23. Henri Gilbert, Rachelle Heim Boissier, Louiza Khati, and Yann Rotella. Generic attack on duplex-based AEAD modes using random function statistics. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology - EUROCRYPT 2023 - 42nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Lyon, France, April 23-27, 2023, Proceedings, Part IV*, volume 14007 of *Lecture Notes in Computer Science*, pages 348–378. Springer, 2023.

JLM14. Philipp Jovanovic, Atul Luykx, and Bart Mennink. Beyond $2^{c/2}$security in sponge-based authenticated encryption modes. In Palash Sarkar and Tetsu Iwata, editors, *Advances in Cryptology - ASIACRYPT 2014 - 20th International Conference on the Theory and Application of Cryptology and Information Security, Kaoshiung, Taiwan, R.O.C., December 7-11, 2014. Proceedings, Part I*, volume 8873 of *Lecture Notes in Computer Science*, pages 85–104. Springer, 2014.

JLM+19. Philipp Jovanovic, Atul Luykx, Bart Mennink, Yu Sasaki, and Kan Yasuda. Beyond conventional security in sponge-based authenticated encryption modes. *J. Cryptol.*, 32(3):895–940, 2019.

Men23. Bart Mennink. Understanding the duplex and its security. *IACR Trans. Symmetric Cryptol.*, 2023(2):1–46, 2023.

MRH04. Ueli M. Maurer, Renato Renner, and Clemens Holenstein. Indifferentiability, Impossibility Results on Reductions, and Applications to the Random Oracle Methodology. In Moni Naor, editor, *Theory of Cryptography, First Theory of Cryptography Conference, TCC 2004, Cambridge, MA, USA, February 19-21, 2004, Proceedings*, volume 2951 of *Lecture Notes in Computer Science*, pages 21–39. Springer, 2004.

MRV15. Bart Mennink, Reza Reyhanitabar, and Damian Vizár. Security of full-state keyed sponge and duplex: Applications to authenticated encryption.

In Tetsu Iwata and Jung Hee Cheon, editors, *Advances in Cryptology - ASIACRYPT 2015 - 21st International Conference on the Theory and Application of Cryptology and Information Security, Auckland, New Zealand, November 29 - December 3, 2015, Proceedings, Part II*, volume 9453 of *Lecture Notes in Computer Science*, pages 465–489. Springer, 2015.