Final-round updates on Romulus

Chun Guo[1], Tetsu Iwata[2], Mustafa Khairallah[3⋆], Kazuhiko Minematsu[4], Thomas Peyrin[5]

[1] Shandong University, China
chun.guo@sdu.edu.cn

[2] Nagoya University, Japan
tetsu.iwata@nagoya-u.jp

[3] Seagate Research, Singapore
mustafa.khairallah@seagate.com

[4] NEC Corporation, Japan
k-minematsu@nec.com

[5] Nanyang Technological University, Singapore
thomas.peyrin@ntu.edu.sg

Webpage: https://romulusae.github.io/romulus/

September 27, 2022

1 New security analysis of Skinny tweakable block cipher

The Skinny twekable block cipher [3] was designed to be secure against related-tweakey attacks, an attack model very generous to the adversary as he can fully control the tweak input. We refer to the original research paper for the extensive security analysis provided by the authors (differential cryptanalysis, linear cryptanalysis, meet-in-the-middle attacks, impossible differential attacks, integral attacks, slide attacks, invariant subspace cryptanalysis, and algebraic attacks). In particular, strong security guarantees for Skinny have been provided with regards to differential and linear cryptanalysis. In addition, since the publication of the cipher in 2016 there has been lots of cryptanalysis or structural analysis (improvement of security bounds) of Skinny by third parties (in order to avoid a very large collection of references, we present below only the current best attacks). This was also further motivated by the organization of cryptanalysis competitions of Skinny by the designers.

We recall that all Romulus versions use internally the Skinny-128-384+ TBC, which is a 40-round reduced version of Skinny-128-384, aiming for 128-bit security. To the best of our knowledge, the cryptanalysis that can attack the highest number of rounds (related-key rectangle attack [7, 26]) can only reach 32 of the 56 rounds of Skinny-128-384, with a very high data/memory/time complexity. For Skinny-128-384+ that aims for 128-bit security, the best distinguisher only reaches 25 rounds out of 40 in the related-key model, the best attack reaches less than 23 rounds in the single-key model. Note that these reduced-round cryptanalysis are very likely inapplicable when placed inside the Romulus operating mode. Even in the hash function setting where the attacker has a lot of control, the best (preimage) cryptanalysis of Romulus-H can only reach 23 rounds [6] (CRYPTO 2021) for a computational complexity of 2^{248} (thus way beyond our security claims). The same authors also described a free-start collision attack on 23 rounds of the compression function of Romulus-H.

All in all, we can conclude that our internal primitive Skinny-128-384+ has a very large security margin (about 40%). The actual Romulus scheme ensures an even larger security margin (> 50%) by limiting the ap-

[⋆] Part of this work was done while the third author was working in Nanyang Technological University, Singapore plicability of these cryptanalysis (fixed TK words, constraints from the operating mode, 2^{64} data limitations, etc.). We believe such a security margin is very important for an encryption algorithm that will probably remain in devices on-field (thus difficult, if not impossible, to update) for decades.

We list below in more details the very recent third-party security analysis of the Skinny tweakable block cipher, our internal primitive. The most explored type of cryptanalysis by third party have been related-key rectangle/boomerang attacks. Hadipour *et al.* [12] (FSE 2022) proposed related-key rectangle attacks on up to 30 rounds of Skinny-128-384 with 2^{361} time and 2^{125} data and Qin *et al.* [25] (FSE 2022) also described related-key rectangle attacks up to 30 rounds of Skinny-128-384 with a slightly lesser complexity of 2^{341} time and 2^{122} data. These related-key rectangle cryptanalysis have later been improved by Dong *et al.* [7] (EUROCRYPT 2022) to 32 rounds of Skinny-128-384, with 2^{355} time and 2^{123} data, using shorter distinguishers but a larger key-recovery part. Very recently, this result has further slightly been improved to 2^{345} time complexity by Song *et al.* [26].

Regardless of the fact that they are in the related-key model, we emphasize that these results are applicable to Skinny-128-384 with a 384-bit key (taking up the entire tweakey material) and not to our internal primitive Skinny-128-384+ that aims for 128-bit security. For example, Hadipour *et al.* [12] can only reach 24 rounds when one TK word of Skinny-128-384 is fixed with 2^{209} time and 2^{125} data, while Qin *et al.* [25] can only reach 25 rounds in this setting with 2^{226} time and 2^{124} data, and Dong *et al.* [7] can only reach 26 rounds with 2^{254} time and 2^{126} data. These lower results still don't apply to Skinny-128-384+ since they aim to attack a 256-bit key.

In terms of related-key distinguisher, Hadipour *et al.* [12] could reach 25 rounds of Skinny-128-384 with probability $2^{-116.6}$ (21 rounds with probability 2^{-114} with one TK word fixed). The related-key boomerang attack proposed by Delaune *et al.* [5] (FSE 2022, received the best paper award) allowed to obtain a related-key boomerang distinguisher on 24 rounds of Skinny-128-384 with 2^{86} time and data (down to 20 rounds when one TK word is fixed).

In the single-key model, the best cryptanalysis on Skinny-128-384 remains the meet-in-the-middle attack from Dong *et al.* [6], breaking 23 rounds with complexity 2^{376} time and 2^{104} data, and more recently the integral attack from Hadipour *et al.* [13], breaking 26 rounds with complexity 2^{344} time and 2^{121} data. We again recall that all these attacks are not applicable to Skinny-128-384+ as their computational cost is way beyond its 128-bit security claims. For comparison, when attacking Skinny-128-256 for a 256-bit security (still way beyond Skinny-128-384+ claims), the authors could only reach 22 rounds with complexity 2^{216}. When attacking Skinny-64-192 for a 128-bit security, the authors could only reach 17 rounds with complexity $2^{116.51}$. We also note a low-data complexity variant of the meet-in-the-middle attack [6] by Hua *et al.* [14], reaching 19 rounds.

Quantum attacks have also been explored with the work of David *et al.* [4] (WCC 2022), who provide quantum impossible differential attacks on 21 rounds of Skinny-128-256 (a very slight complexity improvement over the best non-quantum impossible differential attack on the same primitive). This again doesn't apply to Skinny-128-384+, but it indicates that quantum setting might not be helping the attacker so much with regards to impossible differential attacks on Skinny.

Finally, we mention the recent work from Kuijsters *et al.* [19] that proposes a correlation 1 linear approximation for a certain subset of the tweakey space on the "full 2 rounds" of Skinny. This observation is interesting to better understand the interplay between the internal cipher and the tweakey schedule of Skinny, but not very surprising as Skinny uses a very lightweight tweakey schedule and Sbox (which is of course compensated by the numerous number of rounds). The authors claim that this property on the "full 2 rounds" is a design issue, but we disagree with their statement, as this property has basically no effect on the security of Skinny.

2 New security analysis of Romulus modes

All the members of Romulus are supported by provable security analysis. Specifically, Romulus-N and Romulus-M (the primary and the secondary members) are proved under the standard model, and the proofs are published in ToSC 2020 [17]. For authenticated encryption, provable security in the standard model is an important feature for high security confidence. The current NIST recommendations, AES-GCM and AES-CCM, are also proved under the standard model, namely the pseudorandomness of AES. Among the finalists, only Romulus and GIFT-COFB showed standard model security.

The confidence in the correctness of the provable security results is crucially important. Prof. Jooyoung Lee (KAIST) has undertaken a third-party evaluation of Romulus-N and Romulus-M focusing on the correctness of their security proofs. The report is available at [20]. The report confirms the correctness of the provable security results for Romulus-N and Romulus-M by presenting independent proofs based on a different proof strategy, the H-Coefficient technique. Quote:

> In this evaluation, we proved the security of Romulus-N and Romulus-M; the best attack on any of these modes implies a chosen-plaintext attack (CPA) in the single-key setting against the underlying tweakable block cipher. So unless the tweakable block cipher is broken by CPA adversaries in the single-key setting, Romulus indeed maintains the claimed n-bit security. To evaluate the security of Romulus, with the standard model proof, we can focus on the security evaluation of the underlying primitive. The provable security of Romulus-N and Romulus-M is a clear advantage over any scheme with security proofs in non-standard models.

Several deeper security studies on Romulus-N and Romulus-M appear recently. Habu, Minematsu and Iwata published a paper showing the almost tightness of the provable security bounds of Romulus-M, by presenting matching attacks [11]. This means that the provable security bounds of Romulus-M is optimal for a large class of parameters. Inoue, Guo and Minematsu studied *nonce-misuse resilience* [15] security of Romulus-N, where nonce-misuse resilience is a relaxed security notion from the nonce-misuse resistance [2]. They showed n-bit privacy and $n/2$-bit authenticity in the sense of nonce-misuse resilience. In addition, the authenticity bound has a graceful degradation with respect to nonce repeat/misuse, so if nonce repeat is not frequent, it maintains almost n-bit authenticity. This shows that Romulus-N maintains (almost) n-bit nonce-misuse resilience security, quantitatively comparable to the case of nonce-respecting adversary.

The security of the Romulus-H hash function has been studied and refined [10] from the original MDPH proposal [22]. As a result, Romulus-H maintains the original bit security, $(n - \log n)$-bit indifferentiability security, from the random oracle.

Finally, for our leakage-resilient mode Romulus-T, as promised in the latest specification document, we present the full security proofs of Romulus-T in [9]. The document shows security bounds for stronger notions than conventional privacy and authenticity notions. The results imply our claimed $(n-\log n)$-bit authenticity and privacy for Romulus-T, even in the nonce-misuse resilience setting (i.e., the CCAm$ setting, see [9]). Our proofs for Romulus-T also support the side-channel security claims in the latest specification document, i.e., as long as the side-channel attacker has not recovered the key K, $(n - \log n)$-bit authenticity is kept even with full nonce-misuse, while a birthday-type privacy is easily achieved in nonce-misuse resilience setting.

Importance of long-term security. We emphasize that all the four Romulus modes, Romulus-N, Romulus-M, Romulus-H, and Romulus-T, have a provable security result showing n-bit or $(n - \log n)$-bit security for most of the security notions which corresponds to about 128-bit or 121-bit security when $n = 128$. The NIST standards will be widely used in billions of devices (many of them impossible to update) for the next 30–50 years (or more). We believe that it makes sense to consider the security requirements in 30 years, and the current efficiency comparison takes the provable security results into consideration. Having a proof of security with strong bounds as for Romulus modes greatly reduces the chance of future failures (even of nonce failures), especially for Romulus modes that already have a third party proof.

3 New implementation results

Romulus is particularly efficient on hardware, while generally performing well on software, in particular quite good on very constrained microcontrollers such as 8-bit AVR (see e.g. https://lwc.las3.de/). Below we list some new hardware/software implementation results on Romulus.

Software. An ACISP 2022 paper [1] showed new SIMD implementations of Romulus. They presented two new decompositions of 8-bit S-box of Skinny into 4-bit tables and implemented these tables by vector-permute instructions available in 64-bit ARM or Intel CPUs. They reported speedup gain around 4 to 5 from the previous bitslice/fixslice implementations depending on the platform (also see the Supercop benchmark results[1]. The source available from Github[2]). We remark that their implementations are constant-time.

A SAC 2021 paper showed a new parallel decryption routine, dubbed pincer decryption, for serial MAC or AE modes [21]. The technique ideally doubles the decryption speed on dual-core CPUs or when SIMD instructions are available. Unlike bitslicing, it works for single ciphertext. Pincer decryption is applicable to Romulus, while not to Sponges. The paper reported Romulus-N implementation on ESP32 microcontroller and showed an expected speedup.

Hardware. At NIST LWC workshop 2022, Khairallah and Bhasin presented a configurable combined hardware accelerator for both Romulus-N and Romulus-M [18], showing that both Romulus-N and Romulus-M have lower energy-area product than most candidates except TinyJambu (in terms of unprotected implementations). This is shown in Figure 1. These results make the case not just for Romulus-N, but also for Romulus-M, which provides stronger security guarantees in terms of nonce misuse resistance and security against release of unverified plaintext. They also showed several first-order masked implementations, some of which are shown in Table 1.

Table 1: Examples of synthesis results of different first-order masked implementations of the overall design using Synopsys Design Compiler and TSMC 65nm. The table shows area in GE. All implementations are synthesized for about 2 GHz.

Masking Scheme	Protected Key	Unprotected Key
Domain Oriented Masking	14619.5	13068.47
Consolidated Masking Scheme	15912.7	14372.01
Hardware Private Circuits	18585	17338.75

Steinbauer *et al.* produced a report titled "TVLA On Selected NIST LWC Finalists" [27] which analyzes a different first-order masked implementation or Romulus-N based on hardware private circuits and shows that the implementation passes the TVLA test with 10 million traces, with the implementation being the *smallest among the 5 candidates considered.* It has an area of less than 3,000 LUTs (about half of the largest considered implementation) and requires the least amount of online randomness.

There is a compilation of reports [29] focusing on side-channel security of several finalists, including Romulus. The report, entitled as "On the Side Channel Leakage Assessment of First-Order Masked Romulus", identified only one source of leakage in the same implementation, when varying bits of the nonce. The authors reported not being able to convert this leakage into an attack. We note that this type of leakage is not surprising and, in fact, expected. The nonce is a public parameter. Thus, by fixing all other variables and only changing the nonce, it is normal to observe this type of leakage as the nonce is not expected to be heavily protected.

[1] https://bench.cr.yp.to/results-nistlwc-aead.html

[2] https://github.com/aadomn/skinny

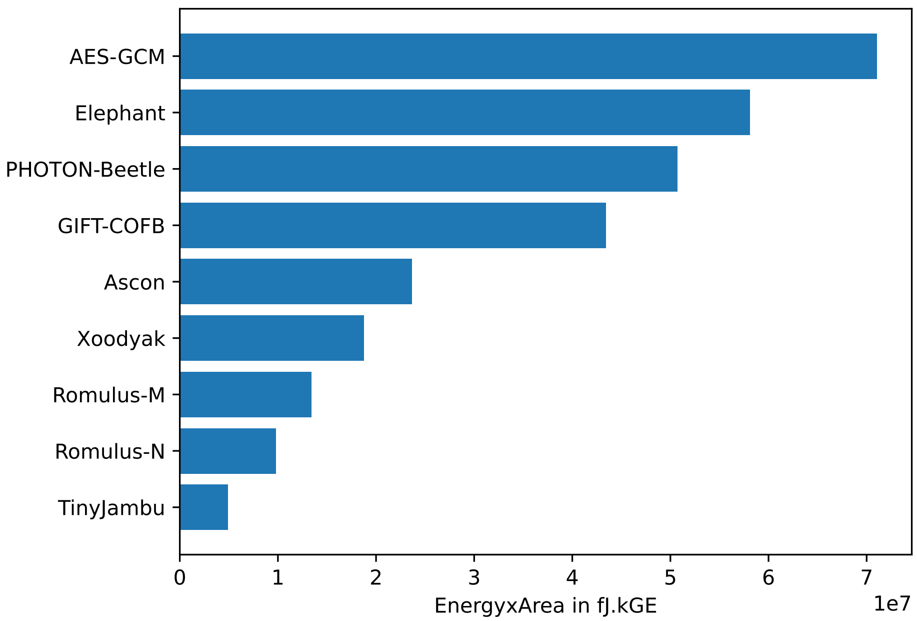

Fig. 1: Comparison of Romulus-N, Romulus-M, AES-GCM and a group of other finalists based on [18]. The synthesis is done using TSMC 65nm.

Remarks on comparison to AES-GCM. Romulus is much more efficient on all hardware metrics (area, energy, power) and for both ASIC and FPGA. Romulus is also much more efficient on very constrained microcontrollers such as 8-bit AVR, and is only slightly better on 32/64-bit microcontrollers, but AES-GCM already performs well on them. One can refer to the benchmarking studies listed at NIST page[3] or our NIST LWC workshop 2022 presentation [8].

4 Other new (related) results

As mentioned in Section 1, numerous cryptographic analyses have been performed on Skinny since the publication, contributing to a significant increase in the reliability of its design. As a result, Skinny has been included in the ISO/IEC standard of tweakable block ciphers (ISO/IEC 18033-7:2022) [16].

At NIST LWC Workshop 2022, Vehamme *et al.* [28] showed an analysis on the finalists from the viewpoints of side-channel leakage resistance. In their analysis, our Romulus-T is classified as one of the schemes that achieve qualitatively the strongest protection under leakage for confidentiality and authenticity.

From a more general perspective, recent studies [23,24] indicate the suitability of tweakable block ciphers for (higher order) masking when operated in an authenticated encryption mode. The basic principle in these studies is that tweaks are easy to protect against side-channel attacks; if they are public no protection is needed, and even for secret tweaks, the protection overhead is small when the tweaeky schedule is linear.

References

1. Adomnicai, A., Minematsu, K., Shigeri, M.: Fast Skinny-128 SIMD Implementations for Sequential Modes of Operation. IACR Cryptol. ePrint Arch. (2022) 578 (To appear at ACISP 2022.).

[3] https://csrc.nist.gov/Projects/lightweight-cryptography/performance-benchmarking

2. Ashur, T., Dunkelman, O., Luykx, A.: Boosting Authenticated Encryption Robustness with Minimal Modifications. In: CRYPTO (3). Volume 10403 of Lecture Notes in Computer Science., Springer (2017) 3–33
3. Beierle, C., Jean, J., Kölbl, S., Leander, G., Moradi, A., Peyrin, T., Sasaki, Y., Sasdrich, P., Sim, S.M.: The SKINNY Family of Block Ciphers and Its Low-Latency Variant MANTIS. In Robshaw, M., Katz, J., eds.: CRYPTO 2016, Part II. Volume 9815 of LNCS., Springer (2016) 123–153
4. David, N., Naya-Plasencia, M., Schrottenloher, A.: Quantum impossible differential attacks: Applications to AES and SKINNY. Cryptology ePrint Archive, Paper 2022/754 (2022) https://eprint.iacr.org/2022/754.
5. Delaune, S., Derbez, P., Vavrille, M.: Catching the Fastest Boomerangs Application to SKINNY. IACR Trans. Symmetric Cryptol. **2020**(4) (2020) 104–129
6. Dong, X., Hua, J., Sun, S., Li, Z., Wang, X., Hu, L.: Meet-in-the-Middle Attacks Revisited: Key-recovery, Collision, and Preimage Attacks. Cryptology ePrint Archive, Paper 2021/427 (2021) https://eprint.iacr.org/2021/427.
7. Dong, X., Qin, L., Sun, S., Wang, X.: Key Guessing Strategies for Linear Key-Schedule Algorithms in Rectangle Attacks. In: EUROCRYPT (3). Volume 13277 of Lecture Notes in Computer Science., Springer (2022) 3–33
8. Guo, C., Iwata, T., Khairallah, M., Minematsu, K., Peyrin, T.: Romulus as NIST LWC Finalist. NIST Lightweight Cryptography Workshop 2022 (2022) https://csrc.nist.gov/csrc/media/Presentations/2022/romulus-as-nist-lwc-finalist/images-media/session-5-peyrin-romulus-as-a-nist-lwc-finalist.pdf.
9. Guo, C., Iwata, T., Khairallah, M., Minematsu, K., Peyrin, T.: Security Proof for Romulus-T. https://romulusae.github.io/romulus/docs/Romulus_T_proof.pdf (2022)
10. Guo, C., Iwata, T., Minematsu, K.: New indifferentiability security proof of MDPH hash function. IET Inf. Secur. **16**(4) (2022) 262–281
11. Habu, M., Minematsu, K., Iwata, T.: Matching Attacks on Romulus-M. IET Inf. Secur., July 2022 (available at https://ietresearch.onlinelibrary.wiley.com/doi/full/10.1049/ise2.12075)
12. Hadipour, H., Bagheri, N., Song, L.: Improved Rectangle Attacks on SKINNY and CRAFT. IACR Trans. Symmetric Cryptol. **2021**(2) (2021) 140–198
13. Hadipour, H., Sadeghi, S., Eichlseder, M.: Finding the Impossible: Automated Search for Full Impossible Differential, Zero-Correlation, and Integral Attacks (Preliminary Version). Cryptology ePrint Archive, Paper 2022/1147 (2022) https://eprint.iacr.org/2022/1147.
14. Hua, J., Liu, T., Cui, Y., Qin, L., Dong, X., Cui, H.: Low-Data Cryptanalysis On SKINNY Block Cipher. The Computer Journal (02 2022) bxab208.
15. Inoue, A., Guo, C., Minematsu, K.: Nonce-Misuse Resilience of Romulus-N and GIFT-COFB. Cryptology ePrint Archive, Paper 2022/1012 (2022) https://eprint.iacr.org/2022/1012.
16. ISO/IEC: 18033-7:2022 – Information security – Encryption algorithms – Part 7: Tweakable block ciphers (2022)
17. Iwata, T., Khairallah, M., Minematsu, K., Peyrin, T.: Duel of the Titans: The Romulus and Remus Families of Lightweight AEAD Algorithms. IACR Trans. Symmetric Cryptol. **2020**(1) (2020) 43–120
18. Khairallah, M., Bhasin, S.: Hardware Implementations of Romulus: Exploring Nonce Misuse Resistance and Boolean Masking. NIST Lightweight Cryptography Workshop 2022 https://csrc.nist.gov/csrc/media/Events/2022/lightweight-cryptography-workshop-2022/documents/papers/hardware-implementations-of-romulus.pdf.
19. Kuijsters, D., Verbakel, D., Daemen, J.: Weak Subtweakeys in SKINNY. Cryptology ePrint Archive, Paper 2022/1042 (2022) https://eprint.iacr.org/2022/1042.
20. Lee, J.: Security Evaluation of Romulus https://romulusae.github.io/romulus/docs/Security_evaluation_Romulus_Jooyoung_Lee.pdf.
21. Minematsu, K., Inoue, A., Moriwaki, K., Shigeri, M., Kubo, H.: Parallel Verification of Serial MAC and AE Modes. In: SAC. Volume 13203 of Lecture Notes in Computer Science., Springer (2021) 200–219
22. Naito, Y.: Optimally Indifferentiable Double-Block-Length Hashing Without Post-processing and with Support for Longer Key Than Single Block. In: LATINCRYPT. Volume 11774 of Lecture Notes in Computer Science., Springer (2019) 65–85
23. Naito, Y., Sasaki, Y., Sugawara, T.: Lightweight Authenticated Encryption Mode Suitable for Threshold Implementation. In: EUROCRYPT (2). Volume 12106 of Lecture Notes in Computer Science., Springer (2020) 705–735
24. Naito, Y., Sasaki, Y., Sugawara, T.: Secret Can Be Public: Low-Memory AEAD Mode for High-Order Masking. In: CRYPTO. Volume 13509 of Lecture Notes in Computer Science., Springer (2022)
25. Qin, L., Dong, X., Wang, X., Jia, K., Liu, Y.: Automated Search Oriented to Key Recovery on Ciphers with Linear Key Schedule Applications to Boomerangs in SKINNY and ForkSkinny. IACR Trans. Symmetric Cryptol. **2021**(2) (2021) 249–291
26. Song, L., Zhang, N., Yang, Q., Shi, D., Zhao, J., Hu, L., Weng, J.: Optimizing Rectangle Attacks: A Unified and Generic Framework for Key Recovery. IACR Cryptol. ePrint Arch. (2022) 723

27. Steinbauer, T., Nagpal, R., Primas, R., Mangard, S.: TVLA On Selected NIST LWC Finalists. https://cryptography.gmu.edu/athena/LWC/Reports/TUGraz/TUGraz_Report_HW_5_candidates_RUB.pdf.
28. Verhamme, C., Cassiers, G., Standaert, F.: Analyzing the Leakage Resistance of the NIST's Lightweight Crypto Standardization Process Finalists. NIST Lightweight Cryptography Workshop 2022 (2022) https://csrc.nist.gov/csrc/media/Events/2022/lightweight-cryptography-workshop-2022/documents/papers/analyzing-the-leakageresistance-of-the-nist-lwc-standardization-process-finalists.pdf.
29. Zhang, X., Wang, T., Cao, P.: Side-Channel Evaluation on Protected Implementations of Several NIST LWC Finalists. https://cryptography.gmu.edu/athena/LWC/Reports/SJTU/SJTU_Report_HW_4_candidates_RUB.pdf.