A Note on ARADI and LLAMA

Roberto Avanzi[1,2], Orr Dunkelman[3] and Shibam Ghosh[3]

[1] Qualcomm Germany GmbH, Munich, Germany
ravanzi@qti.qualcomm.com
[2] Caesarea Rothschild Institute, University of Haifa, Haifa, Israel
roberto.avanzi@gmail.com
[3] Computer Science Department, University of Haifa, Haifa, Israel
orrd@cs.haifa.ac.il, sghosh03@campus.haifa.ac.il

Abstract. Recently, the NSA has proposed a block cipher called ARADI and a mode of operation called LLAMA for memory encryption applications. In this note, we comment on this proposal, on its suitability for the intended application, and describe an attack on LLAMA that breaks confidentiality of ciphertext and allows a straightforward forgery attack breaking integrity of ciphertext (*INT-CTXT*) using a related-*Initialization Vector* (IV) attack. Both attacks have negligible complexity.

Keywords: Block Ciphers · Lightweight Cryptography · Modes of Operation · Memory Encryption

1 Introduction

Greene, Motley and Weeks of the NSA have introduced the block cipher ARADI and the mode of operation LLAMA for the use case of confidentiality and integrity of computer memory [GMW24]. This is an active area of research in both academia and industry, where encryption can be performed by standard ciphers like the AES [DR02], PRESENT [BKL+07], or ad-hoc lightweight ciphers such as PRINCE [BCG+12], QARMA [Ava17] or QARMAv2 [ABD+23].

We first note that ARADI, according to its own authors (cf. Section 3.5 of [GMW24]), does not offer a major area advantage over the AES. A significant performance advantage over the AES, in other words an appreciably lower latency, is achieved only with a much larger area.

The area-latency product is between 1.3 and 3.7 times lower than that of the AES. These values are similar to those obtained for the tweakable block cipher QARMA [Ava17] (for instance, the area-latency product of $\texttt{QARMA}_9$-128-σ_1 is about 3.33 times better the AES's) and therefore the new design does not obviously seem better than the current state of the art. Also, since ARADI, like the AES, is not a tweakable design, either latency, area, or both may degrade by a factor up to 2 depending on the application. This suggests that the performance comparison of [GMW24] should be improved.

LLAMA's Suitability For Memory Encryption LLAMA is a counter-based mode of encryption, whereby a keystream is generated by encrypting successive values of the concatenation of an IV and a counter, and XORing the keystream to the plaintext to obtain the ciphertext. Such an encryption mode suffers from a few limitations in the context of computer memory:

1. If there is a ciphertext leakage channel while the machine is running, the method requires fresh IVs at each write to the same location, otherwise the XOR of two plaintexts is immediately recovered. This implies that the IVs must be stored somewhere

in memory, which either takes space on the chip or reduces both memory availability and bandwidth. In fact, in [AMS+22] it is shown that the memory requirements to store the counter and the associated memory accesses can have a significant impact on overall system performance, to the point that in many cases direct encryption methods are preferred and replace previous keystream based methods (cf. the involution from SGX [Gue16] to Scalable SGX [JMSS20]).

2. If ciphertext leakage can occur only through warm- and cold-boot attacks [HSH+09], then the security model is weaker, but one can use a fixed IV per cache line.

Attacking `LLAMA` Despite years of research in symmetric-key design, it is customary that designers add their security analysis. The authors of [GMW24] do not offer any security claim that can be verified. A user of `ARADI` does not have knowledge about the claimed safety margins of the design. While other NSA designs such as `Simon` and `Speck` [BSS+13] were also released without those claims, we note that `LLAMA` is proposed without any security claims nor proofs, as common in modes of operation's research. This leads us to closely examine `LLAMA`, and realize its lack of concrete security claims and proof is due to a simple nonce respecting related-IV attack that breaks the integrity and confidentiality of the ciphertexts.

Unlike the general purpose `CTR` mode that treats different IV lengths differently, `LLAMA` does not. The counter mode encryption is performed by concatenating an ℓ-bit IV with a $(128-\ell)$-bit counter, and encrypts the result to obtain a 128-bit block of the keystream. However, the IV length is not fixed, so one could, in theory, have two IVs i_0 and $i_1 = i_0\|0$, of lengths, say, 96 and 97 bits, and for the first values of a 31-bit counter c, if $c' = 0\|c$, then $i_0\|c' = i_1\|c$ and therefore the same keystream is generated. This immediately allows breaking the *PRF* security (by decrypting under one IV to allow decrypting under the second). Also, we can mount a straightforward *INT-CTXT* attack: Suppose that we obtain a ciphertext-tag pair (C, T) for some message M and the IV i_0, then (C, T) is a valid forgery under the IV i_1.

While the above attack is probably not a concern for memory encryption, where the IV size is most likely fixed, users of `LLAMA` should be aware of this weakness. This raises some suspicion about the soundness of the whole design.

2 Conclusion

In this note we have discussed a few shortcomings of the newly proposed `ARADI` and `LLAMA`, suggested by [GMW24]. We have shown that the construction is probably not competitive for memory encryption, and suggested that the `LLAMA` mode of operation may not offer enough security in different context. To conclude, the proposal of `ARADI` and `LLAMA` lacks design rationale, any concrete security claims and security proofs.

References

[ABD+23] Roberto Avanzi, Subhadeep Banik, Orr Dunkelman, Maria Eichlseder, Shibam Ghosh, Marcel Nageler, and Francesco Regazzoni. The QARMAv2 Family of Tweakable Block Ciphers. *IACR Transactions on Symmetric Cryptology*, (3):25–73, Sep. 2023. doi:10.46586/tosc.v2023.i3.25-73. Cited on page 1.

[AMS+22] Roberto Avanzi, Ionut Mihalcea, David Schall, Héctor Montaner, and Andreas Sandberg. Hardware-Supported Cryptographic Protection of Random Access Memory. Cryptology ePrint Archive, Paper 2022/1472, 2022. https://eprint.iacr.org/2022/1472. Available from: https://eprint.iacr.org/2022/1472. Cited on page 2.

[Ava17] Roberto Avanzi. The QARMA Block Cipher Family — Almost MDS Matrices over Rings with Zero Divisors, Nearly Symmetric Even-Mansour Constructions with Non-Involutory Central Rounds, and Search Heuristics for Low-Latency S-Boxes. *IACR Trans. on Symmetric Cryptology*, 2017(1):4–44, 2017. doi:10.13154/tosc.v2017.i1.4-44. Cited on page 1.

[BCG+12] Julia Borghoff, Anne Canteaut, Tim Güneysu, Elif Bilge Kavun, Miroslav Knezevic, Lars R. Knudsen, Gregor Leander, Ventzislav Nikov, Christof Paar, Christian Rechberger, Peter Rombouts, Søren S. Thomsen, and Tolga Yalçin. PRINCE — A Low-Latency Block Cipher for Pervasive Computing Applications — Extended Abstract. In Xiaoyun Wang and Kazue Sako, editors, *Advances in Cryptology — ASIACRYPT 2012 — 18th International Conference on the Theory and Application of Cryptology and Information Security, Beijing, China, December 2-6, 2012. Proceedings*, volume 7658 of *Lecture Notes in Computer Science*, pages 208–225. Springer, 2012. doi:10.1007/978-3-642-34961-4_14. Cited on page 1.

[BKL+07] Andrey Bogdanov, Lars R. Knudsen, Gregor Leander, Christof Paar, Axel Poschmann, Matthew J. B. Robshaw, Yannick Seurin, and C. Vikkelsoe. PRESENT: An Ultra-Lightweight Block Cipher. In Pascal Paillier and Ingrid Verbauwhede, editors, *Cryptographic Hardware and Embedded Systems — CHES 2007, 9th International Workshop, Vienna, Austria, September 10-13, 2007, Proceedings*, volume 4727 of *Lecture Notes in Computer Science*, pages 450–466. Springer, 2007. Available from: https://doi.org/10.1007/978-3-540-74735-2, doi:10.1007/978-3-540-74735-2_31. Cited on page 1.

[BSS+13] Ray Beaulieu, Douglas Shors, Jason Smith, Stefan Treatman-Clark, Bryan Weeks, and Louis Wingers. The SIMON and SPECK Families of Lightweight Block Ciphers. Cryptology ePrint Archive, Report 2013/404, 2013. Available from: http://eprint.iacr.org/2013/404. Cited on page 2.

[DR02] Joan Daemen and Vincent Rijmen. *The Design of Rijndael: AES — The Advanced Encryption Standard*. Information Security and Cryptography. Springer, 2002. doi:10.1007/978-3-662-04722-4. Cited on page 1.

[GMW24] Patricia Greene, Mark Motley, and Bryan Weeks. ARADI and LLAMA: Low-Latency Cryptography for Memory Encryption. Cryptology ePrint Archive, Paper 2024/1240, 2024. https://eprint.iacr.org/2024/1240. Available from: https://eprint.iacr.org/2024/1240. Cited on pages 1 and 2.

[Gue16] Shay Gueron. A Memory Encryption Engine Suitable for General Purpose Processors. *IACR Cryptol. ePrint Arch.*, 2016. Available from: http://eprint.iacr.org/2016/204. Cited on page 2.

[HSH+09] J. Alex Halderman, Seth D. Schoen, Nadia Heninger, William Clarkson, William Paul, Joseph A. Calandrino, Ariel J. Feldman, Jacob Appelbaum, and Edward W. Felten. Lest we remember: cold-boot attacks on encryption keys. *Commun. ACM*, 52(5):91–98, 2009. doi:10.1145/1506409.1506429. Cited on page 2.

[JMSS20] Simon Johnson, Raghunandan Makaram, Amy Santoni, and Vinnie Scarlata. Supporting Intel® SGX on Multi-Socket Platforms, August 2020. Technical Report. Cited on page 2.